

A Generalization of the Łoś-Tarski Preservation Theorem

Submitted in partial fulfillment of
the requirements of the degree of

Doctor of Philosophy

by

Abhisekh Sankaran
(Roll No. 07405001)

Under the guidance of
Prof. Supratik Chakraborty
and
Prof. Bharat Adsul



Department of Computer Science and Engineering
Indian Institute of Technology Bombay

August 2016

To my parents

Thesis Approval

The thesis entitled

A Generalization of the Łoś-Tarski Preservation Theorem

by

Abhisekh Sankaran

(Roll No. 07405001)

is approved for the degree of

Doctor of Philosophy

Prof. Anuj Dawar

University of Cambridge, UK

External Examiner

Prof. Ajit A. Diwan

Dept. of CSE, IIT Bombay

Internal Examiner

Prof. Supratik Chakraborty

Dept. of CSE, IIT Bombay

Supervisor

Prof. Bharat Adsul

Dept. of CSE, IIT Bombay

Co-Supervisor

Prof. Harish K. Pillai

Dept. of EE, IIT Bombay

Chairman

Date: _____

Place: _____

Declaration

I declare that this written submission represents my ideas in my own words and where others' ideas or words have been included, I have adequately cited and referenced the original sources. I also declare that I have adhered to all principles of academic honesty and integrity and have not misrepresented or fabricated or falsified any idea/data/fact/source in my submission. I understand that any violation of the above will be cause for disciplinary action by the Institute and can also evoke penal action from the sources which have thus not been properly cited or from whom proper permission has not been taken when needed.

(Signature of student)

Abhisekh Sankaran

(Name of student)

Roll No. **07405001**

Date: _____

Abstract

Preservation theorems are amongst the earliest areas of study in classical model theory. One of the first preservation theorems to be proven is the Łoś-Tarski theorem that provides over arbitrary structures and for arbitrary finite vocabularies, semantic characterizations of the \forall^* and \exists^* prefix classes of first order logic (FO) sentences, via the properties of preservation under substructures and preservation under extensions respectively. In the classical model theory part of this thesis, we present new parameterized preservation properties that provide for each natural number k , semantic characterizations of the $\exists^k \forall^*$ and $\forall^k \exists^*$ prefix classes of FO sentences, over the class of all structures and for arbitrary finite vocabularies. These properties, that we call *preservation under substructures modulo k -cruxes* and *preservation under k -ary covered extensions* respectively, correspond exactly to the properties of preservation under substructures and preservation under extensions, when k equals 0. As a consequence, we get a parameterized generalization of the Łoś-Tarski theorem for sentences, in both its substructural and extensional forms. We call our characterizations collectively the *generalized Łoś-Tarski theorem for sentences at level k* , abbreviated $\text{GLT}(k)$. To the best of our knowledge, $\text{GLT}(k)$ is the first to relate *counts* of quantifiers appearing in the sentences of the Σ_2^0 and Π_2^0 prefix classes of FO, to natural quantitative properties of models, and hence provides new semantic characterizations of these sentences. We generalize $\text{GLT}(k)$ to theories, by showing that theories that are preserved under k -ary covered extensions are characterized by theories of $\forall^k \exists^*$ sentences, and theories that are preserved under substructures modulo k -cruxes, are equivalent, under a well-motivated model-theoretic hypothesis, to theories of $\exists^k \forall^*$ sentences. We also present natural variants of our preservation properties in which, instead of natural numbers k , we consider infinite cardinals λ , and show that these variants provide new semantic characterizations of Σ_2^0 and Π_2^0 theories. In contrast to existing preservation properties in the literature that characterize Σ_2^0 and Π_2^0 sentences, our preservation properties are combinatorial and finitary in nature, and stay non-trivial over finite structures as well. Hence, in the finite model theory part of the thesis, we investigate

GLT(k) over finite structures. Like most preservation theorems, GLT(k) fails over the class of all finite structures. To “recover” GLT(k), we identify a new logic based combinatorial property of classes \mathcal{S} of finite structures, that we call the *\mathcal{L} -equivalent bounded substructure property*, abbreviated \mathcal{L} -EBSP(\mathcal{S}, k), where \mathcal{L} is either FO or MSO. We show that \mathcal{L} -EBSP(\mathcal{S}, k) entails GLT(k), and even an effective version of the latter, under suitable “computability” assumptions. A variety of classes of finite structures of interest in computer science turn out to satisfy \mathcal{L} -EBSP(\mathcal{S}, k), and the just mentioned computability assumptions as well, whereby all of these classes satisfy an effective version of GLT(k). Examples include the classes of words, trees (unordered, ordered or ranked), nested words, cographs, graph classes of bounded tree-depth, graph classes of bounded shrub-depth and n -partite cographs. These classes were earlier not known to even satisfy the Łoś-Tarski theorem. All of the aforesaid classes have received significant attention due to their excellent logical and algorithmic properties, and moreover, many of these are recently defined (in the last 10 years). We go further to give ways to construct new classes of structures satisfying \mathcal{L} -EBSP(\cdot, \cdot) by showing the closure of the latter under set-theoretic operations and special kinds of translation schemes. As a consequence, we get that \mathcal{L} -EBSP(\cdot, k) is closed under unary operations like complementation, transpose and the line-graph operation, and binary “sum-like” operations like disjoint union and join, while FO-EBSP(\cdot, \cdot) is additionally closed under “product-like” operations like the cartesian, tensor, lexicographic and strong products. On studying the \mathcal{L} -EBSP(\cdot, k) property further, it turns out that any class of structures that is well-quasi-ordered under embedding satisfies \mathcal{L} -EBSP($\cdot, 0$), that \mathcal{L} -EBSP(\cdot, k) classes, under the aforementioned computability assumptions, admit decidability of the satisfiability problem for \mathcal{L} , and that \mathcal{L} -EBSP(\cdot, k) entails the homomorphism preservation theorem. Finally, we find it worth mentioning that \mathcal{L} -EBSP(\mathcal{S}, k) has a remarkably close resemblance to the classical downward Löwenheim-Skolem property, and can very well be regarded as a finitary analogue of the latter. It is pleasantly surprising that while the downward Löwenheim-Skolem property is by itself meaningless over finite structures, a natural finitary analogue of it is satisfied by a wide spectrum of classes of finite structures, that are of interest and importance in computer science.

In summary, the properties introduced in this thesis are interesting in both the classical and finite model theory contexts, and yield in both these contexts, a new parameterized generalization of the Łoś-Tarski preservation theorem.

Contents

Abstract	vii
List of Figures	xiii
List of publications	xv
1 Introduction	1
1.1 Our results	5
1.1.1 Results in the classical model theory context	7
1.1.2 Results in the finite model theory context	9
1.2 Organization of the thesis	13
I Classical Model Theory	15
2 Background and preliminaries	17
2.1 Syntax and semantics of FO	17
2.2 Σ_n^0 and Π_n^0 formulae	21
2.3 Notions concerning structures	22
2.4 Types and μ -saturation	24
2.5 Two classical preservation properties	25
2.5.1 The Łoś-Tarski preservation theorem	26
3 New parameterized preservation properties	29
3.1 Preservation under substructures modulo k -cruxes	29
3.2 Preservation under k -ary covered extensions	31

4	Characterizations: the case of sentences	35
4.1	The generalized Łoś-Tarski theorem for sentences – $\text{GLT}(k)$	35
4.1.1	Proof of $\text{GLT}(k)$ using λ -saturated structures	36
4.1.2	Proof of $\text{GLT}(k)$ using ascending chains of structures	37
4.2	Variants of our properties and their characterizations	40
4.2.1	Applications: new proofs of inexpressibility results in FO	42
4.3	An uncomputability result	42
5	Characterizations: the case of theories	45
5.1	Characterizations of the extensional properties	45
5.2	Characterizations of the substructural properties	47
6	Directions for future work	63
II	Finite Model Theory	67
7	Background and preliminaries	69
7.1	Syntax and semantics of MSO	69
7.2	Adaptations of classical model theory concepts to finite model theory	71
7.3	\mathcal{L} -Ehrenfeucht-Fräïssé games	72
7.4	Translation schemes	74
8	The need to investigate new classes of finite structures for $\text{GLT}(k)$	79
8.1	Failure of $\text{GLT}(k)$ over all finite structures	80
8.2	Failure of $\text{GLT}(k)$ over classes well-behaved w.r.t. the Łoś-Tarski theorem	83
9	The \mathcal{L}-Equivalent Bounded Substructure Property – $\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$	87
9.1	$\mathcal{L}\text{-EBSP}(\cdot, k)$ entails $\text{GLT}(k)$	88
9.2	$\mathcal{L}\text{-EBSP}(\cdot, k)$ and the downward Löwenheim-Skolem property	90
9.3	A sufficient condition for $\mathcal{L}\text{-EBSP}(\cdot, k)$	92
10	Classes satisfying $\mathcal{L}\text{-EBSP}(\cdot, k)$	95
10.1	An abstract result concerning tree representations	95
10.2	Words, trees and nested words	100

10.3	n -partite cographs	109
10.4	Closure properties of \mathcal{L} -EBSP(\cdot, \cdot)	114
10.4.1	Closure under set-theoretic operations	114
10.4.2	Closure under operations implemented using translation schemes	115
10.4.3	Closure under regular operation-tree languages	122
11	Additional studies on \mathcal{L}-EBSP(\cdot, k)	127
11.1	\mathcal{L} -EBSP(\mathcal{S}, k) and the decidability of \mathcal{L} -Th(\mathcal{S})	127
11.2	\mathcal{L} -EBSP(\cdot, k) and well-quasi-ordering under embedding	128
11.3	\mathcal{L} -EBSP(\cdot, k) and the homomorphism preservation theorem	131
12	Directions for future work	137
13	A summary of our contributions	141
	Bibliography	145
	Acknowledgments	153
	Index to symbols	159
	Index	161

List of Figures

4.1	Characterizations of $PSC(k)$, $PCE(k)$, $PSC(\lambda)$ and $PCE(\lambda)$ sentences . . .	41
5.1	Characterizations of $PCE(k)$ and $PCE(\lambda)$ theories	46
5.2	(Partial) characterizations of $PSC(k)$ and $PSC(\lambda)$ theories	53
5.3	(Partial) conditional characterizations of $PSC(k)$ and $PSC(\lambda)$ theories	61
9.1	\mathcal{L} -EBSP(\mathcal{S}, k) as a finitary analogue of \mathcal{L} -DLSP(\mathcal{S}, κ)	92
10.1	Nested Σ -word as a tree over $\Sigma \cup (\Sigma \times \Sigma) \cup \{\circ\}$	107
10.2	Tree over $\Sigma \cup (\Sigma \times \Sigma) \cup \{\circ\}$ as a nested Σ -word	107
10.3	n -partite cograph G and an n -partite cotree representation t of G	110

List of Publications

I. Classical model theory

1. Abhisekh Sankaran, Bharat Adsul, and Supratik Chakraborty. A generalization of the Łoś-Tarski preservation theorem. *Annals of Pure and Applied Logic*, 167(3):189 - 210, 2016.

II. Finite model theory

1. Abhisekh Sankaran, Bharat Adsul, and Supratik Chakraborty. A generalization of the Łoś-Tarski preservation theorem over classes of finite structures. In *Proceedings of the 39th International Symposium on Mathematical Foundations of Computer Science, MFCS 2014, Budapest, Hungary, August 25-29, 2014, Part I*, pages 474 - 485, 2014.
2. A. Sankaran, B. Adsul, V. Madan, P. Kamath, and S. Chakraborty. Preservation under substructures modulo bounded cores. In *Proceedings of the 19th International Workshop on Logic, Language, Information and Computation, WoLLIC 2012, Buenos Aires, Argentina, September 3-6, 2012*, pages 291 - 305, 2012.

Chapter 1

Introduction

Classical model theory is a subject within mathematical logic, that studies the relationship between a formal language and its interpretations, also called structures or models [12, 40]. The most well-studied formal language in classical model theory is *first order logic* (henceforth called FO), a language that is built up from predicates, functions and constant symbols using boolean connectives, and existential and universal quantifications. Classical model theory largely studies the correspondence between the *syntax* of a description in FO with the *semantics* of the description, where the latter is the class of all structures that satisfy the description [82]. Amongst the earliest areas of study in classical model theory, is a class of results called *preservation theorems*. A preservation theorem identifies syntactic features that capture a *preservation property*, which is a special kind of semantics that defines classes of arbitrary structures (that is, structures that could be finite or infinite) that are closed or *preserved* under some model-theoretic operation. For instance, the class of all cliques (graphs in which any two vertices are adjacent) is preserved under the operation of taking substructures (which are induced subgraphs in this context). The class of all cliques is defined by the FO sentence that says “for all (vertices) x and for all (vertices) y , (there is an) edge between x and y ”. The latter is a description in FO having the special syntactic feature that it contains only universal quantifications and no existential quantifications. One of the earliest preservation theorems of classical model theory, the Łoś-Tarski theorem, proven by Jerzy Łoś and Alfred Tarski in 1954-55 [40], says that the aforementioned syntactic feature is indeed *expressively complete* for the semantics of preservation under substructures. In other words, a class of arbitrary structures that is defined by an FO sentence, is preserved under substructures if, and only if, it is definable by a universal sentence, the latter being an FO sentence in which only universal quantifications appear (Theorem

3.2.2 in [12]). In “dual” form, the Łoś-Tarski theorem states that a class of arbitrary structures that is defined by an FO sentence, is preserved under extensions if, and only if, it is definable by an existential sentence which is an FO sentence that uses only existential quantifications. The theorem extends to theories (sets of sentences) as well: a class of arbitrary structures that is defined by an FO theory, is preserved under substructures (respectively, extensions) if, and only if, it is definable by a theory of universal (respectively, existential) sentences. Historically speaking, the study of preservation theorems began with Marczewski asking in 1951, which FO definable classes of structures are preserved under surjective homomorphisms [41]. This question triggered off an extensive study of preservation theorems in which a variety of model-theoretic operations like substructures, extensions, homomorphisms, unions of chains, direct products, reduced products, etc. were taken up and preservation theorems for these operations were proven.

The Łoś-Tarski theorem holds a special place amongst preservation theorems, for its significance from at least two points of view: historical and technical. From the historical point of view, the theorem was amongst the earliest applications of the *compactness theorem* (Gödel 1930, Mal'tsev 1936) [40], a result that is now regarded as one of the pillars of classical model theory. Further, the method of proof of the Łoś-Tarski theorem lent itself to adaptations that enabled proving the various other preservation theorems mentioned above. This extensive research into preservation theorems from the '50s to the '70s (indeed these theorems were subsequently also proven for extensions of FO, like infinitary logics [47]) contributed much to the development of classical model theory [41]. From the technical point of view, the property of preservation under substructures that the Łoś-Tarski theorem characterizes, has been studied substantially in the literature of various mathematical disciplines, under the name of *hereditarity*. A property is hereditary if for any structure satisfying the property, any substructure of it also satisfies the property. Hereditary properties or their variants have been of significant interest in topology, set theory, graph theory and poset theory, to name a few areas. In more detail, in topology, the notions of second countability and metrisability are hereditary, while those of sequentiality and Hausdorff compactness are what are called *weakly hereditary* [48]. In set theory, the notions of hereditary sets, hereditarily finite sets and hereditarily countable sets are all hereditary properties [51]. Various classes of graphs of interest in graph theory are hereditary; examples include cliques, forests, n -partite graphs, planar graphs, graphs of bounded degree, graphs that exclude any fixed finite set of graphs as subgraphs or induced subgraphs [19]. In

poset theory, a landmark result in the sub-area of well-quasi-orders [59], namely the Robertson-Seymour theorem [65], characterizes a variant of hereditariness, called *minor-hereditariness*. The Łoś-Tarski theorem studies hereditariness from the point of view of logic, and specifically, provides a syntactic characterization of hereditary classes of structures that are FO definable.

While a preservation theorem can be seen as providing a syntactic characterization of a preservation property, the same theorem, flipped around, can also be seen as providing a semantic characterization (and furthermore, via a preservation property) of a syntactic class of FO theories. Thus, the Łoś-Tarski theorem provides semantic characterizations of existential and universal theories, in terms of preservation under extensions and preservation under substructures respectively. Existential and universal theories are equivalent respectively to what are known in the literature as Σ_1^0 and Π_1^0 theories. For $n \geq 1$, a Σ_n^0 theory is a set of Σ_n^0 sentences, where a Σ_n^0 sentence is a *prenex* FO sentence in which from left to right, there is a *quantifier prefix* consisting of n blocks of quantifiers (equivalently, $n - 1$ alternations of quantifiers) beginning with a block of existential quantifiers, followed by a quantifier-free formula. Likewise a Π_n^0 theory is a set of Π_n^0 sentences, where a Π_n^0 sentence is a prenex FO sentence in which from left to right, there is a quantifier prefix consisting of n blocks of quantifiers beginning with a block of universal quantifiers, followed by a quantifier-free formula. The Łoś-Tarski theorem provides semantic characterizations of Σ_1^0 and Π_1^0 theories. For Σ_n^0 sentences and Π_n^0 theories for $n \geq 2$, semantic characterizations were proven using preservation properties defined in terms of *ascending chains* and *descending chains* (Theorem 5.2.8 in [12]). Finally in 1960, Keisler proved the *n-sandwich theorem* [45] that provides a characterization of Σ_n^0 and Π_n^0 sentences and theories, for each $n \geq 1$, using preservation properties defined uniformly in terms of the notion of *n-sandwiches*. It is important to note that all of the characterizations mentioned above are over arbitrary structures, and make important use of the presence of infinite structures.

In 1973, Fagin proved a remarkable syntax-semantics correspondence over *finite* structures. He showed that an isomorphism-closed class of finite structures has the (algorithmic) semantic property of being in the complexity class **NP** (Non-deterministic Polynomial time) if, and only if, it is definable in an extension of FO called *existential second order logic* (Theorem 3.2.4 in [34]). This result gave birth to the area of *finite model theory*, whose aims are similar to classical model theory (i.e. study of the expressive power of formal languages) but now the structures under consideration are only finite. Finite model theory [34, 54] is closely connected with computer science since many disciplines within the latter use formal languages, such as

programming languages, database query languages or specification languages, and further the structures that arise in these disciplines are often finite, such as data structures, databases or program models respectively. It is natural to ask if the results and techniques of classical model theory can be carried over to the finite. Unfortunately, it turns out that many important results and methods of classical model theory fail in the context of finite structures. The most stark failure is that of the compactness theorem, whereby, all proofs based on the compactness theorem – indeed this includes the proofs of almost every preservation theorem – fail when restricted to only finite structures. But worse still, the statements of most preservation theorems fail too. The Łoś-Tarski theorem fails in the finite; Tait [81] showed there is an FO sentence that is preserved under substructures over the class of all finite structures, but that is not equivalent over this class, to any universal sentence. The other preservation theorems from the classical model theory literature mentioned earlier, namely those characterizing Σ_n^0 and Π_n^0 sentences/theories for $n \geq 2$, fail in the finite too; this is simply because the characterizing notions become trivial over finite structures. Two rare theorems that survive passage into the finite are the modal characterization theorem and the homomorphism preservation theorem – the former was shown by Rosen [67], and the latter was a striking result due to Rossman [70], that settled a long standing open problem in finite model theory concerning the status of this theorem in the finite. But then, these results are exceptions. (See [1, 2, 3, 33, 37, 66, 69, 80] for more on the investigations of results from classical model theory in the context of all finite structures. See [68] for an excellent survey of these.)

To “recover” classical preservation theorems in the finite model theory setting, recent research [6, 7, 15, 16, 21, 38] in the last ten years, has focussed attention on studying these theorems over “well-behaved” classes of finite structures. In particular, Atserias, Dawar and Grohe showed in [7] that under suitable closure assumptions, classes of structures that are acyclic or of bounded degree admit the Łoś-Tarski theorem for sentences. Likewise, the class of all structures of treewidth at most k also admits the Łoś-Tarski theorem, for each natural number k . These classes of structures are well-behaved in the sense that they have proved especially important in modern graph structure theory as also from an algorithmic point of view [20]. A classic result from graph structure theory states that a minor-hereditary class of graphs has bounded treewidth if, and only if, the class has a finite set of forbidden minors that includes a planar graph [64]. From an algorithmic point of view, many computational problems that are otherwise intractable (such as 3-colorability), become tractable when restricted to structures of bounded treewidth [14].

Likewise, over structures of bounded degree, many problems that are polynomial time solvable in general (such as checking if a graph is triangle-free) become solvable in linear time [77]. Atserias, Dawar and Kolatis showed that the homomorphism preservation theorem also holds over the aforesaid classes of structures [6]. (Note that this theorem being true over all finite structures does not imply that it would be true over subclasses of finite structures; restricting attention to a subclass weakens both the hypothesis and the consequent of the statement of the theorem). Subsequently, Harwath, Heimberg and Schweikardt [38] studied the bounds for an effective version of the Łoś-Tarski theorem and the homomorphism preservation theorem over bounded degree structures. In [21], Duris showed that the Łoś-Tarski theorem holds for structures that are acyclic in a more general sense. All of the classes of structures mentioned above are thus “well-behaved” from the model-theoretic point of view as well (in that, these classes admit theorems from classical model theory). The investigation of such model-theoretic well-behavedness is a current and active area of research in finite model theory.

1.1 Our results

The properties in the classical model theory literature that characterize Σ_n^0 and Π_n^0 sentences or theories, characterize these syntactic classes “as a whole”. None of these characterize Σ_n^0 and Π_n^0 sentences/theories in which for some given block, the number of quantifiers in that block is *fixed* to a given natural number k . Further, all of these properties are in terms of notions that are “infinitary”, i.e. notions that are non-trivial only when arbitrary (i.e. finite and infinite) structures are considered, and that become trivial when restricted to only finite structures. Given the active interest in preservation theorems in the finite model theory context, none of the properties mentioned above can be used to characterize Σ_n^0 and Π_n^0 sentences in the finite, for $n \geq 2$. Further, the preservation theorems that have been investigated over well-behaved classes in the finite model theory literature, namely the Łoś-Tarski theorem and the homomorphism preservation theorem, are those that characterize only Σ_1^0 or Π_1^0 sentences, or subclasses of these.

The observations above raise the following two natural questions:

- Q1. Are there properties that semantically characterize, over arbitrary structures, Σ_n^0 and Π_n^0 sentences/theories in which the number of quantifiers appearing in a given block(s) is fixed to a given natural number(s)?

Q2. Are there properties that semantically characterize, over classes of finite structures, Σ_n^0 and Π_n^0 sentences in which the number of quantifiers appearing in a given block(s) is fixed to a given natural number(s)? If so, what are these classes?

In this thesis, we consider the case when $n = 2$, and present our partial results towards addressing the above questions. Specifically, for the case of Σ_2^0 and Π_2^0 sentences, in which the number of quantifiers in the *leading block* is fixed to a given natural number, we identify preservation properties that uniformly answer both **Q1** and (the first part of) **Q2** in the affirmative. In other words, we present quantitative dual parameterized preservation properties that are finitary and combinatorial, and that characterize over arbitrary structures and over a variety of interesting classes of finite structures, Σ_2^0 and Π_2^0 sentences whose quantifier prefixes are respectively of the form $\exists^k \forall^*$ or $\forall^k \exists^*$ (i.e. k quantifiers in the first quantifier block followed by zero or more quantifiers in the second quantifier block). Our properties, that we call *preservation under substructures modulo k -cruxes* and *preservation under k -ary covered extensions* are exactly the classical properties of preservation under substructures and preservation under extensions for the case of $k = 0$. Whereby, our characterizations of $\exists^k \forall^*$ and $\forall^k \exists^*$ sentences yield the Łoś-Tarski theorem for sentences for the case of $k = 0$. We hence call our characterizations collectively as the *generalized Łoś-Tarski theorem for sentences at level k* , and denote it as $\text{GLT}(k)$. To the best of our knowledge, our characterizations are the first to relate natural quantitative properties of models of sentences in a semantic class to counts of leading quantifiers in equivalent Σ_2^0 or Π_2^0 sentences. Before we present our results in more detail and provide our answer to the second part of **Q2**, we briefly describe the importance of the Σ_2^0 and Π_2^0 classes of sentences.

After Hilbert posed the *Entscheidungsproblem* in 1928, namely the problem of deciding if a given FO sentence is satisfiable, abbreviated the SAT problem, one of the first classes of FO sentences for which SAT was shown to be decidable, was the Σ_2^0 class. This was shown by Bernays and Schönfinkel for Σ_2^0 sentences without equality, and later extended to full Σ_2^0 by Ramsey [63] (on a historical note: it was in showing this result that Ramsey proved the famous *Ramsey's theorem*). In a subsequent extensive research of about 70 years on the SAT problem for prefix classes, it was shown [10] that Σ_2^0 is indeed one of the *maximal* prefix classes for which the SAT problem is decidable. Interestingly, on the other hand, various subclasses of Π_2^0 class turn out to be *minimal* prefix classes for which the SAT problem is undecidable; for instance, the class of Π_2^0 sentences with only two universal quantifiers and over a vocabulary

containing just one binary relation symbol, is undecidable for SAT, when equality is allowed. With the growth of parameterized complexity theory [20], it became interesting to study the computational complexity of the satisfiability problem for the Σ_2^0 class, in terms of *counts of quantifiers* as parameters. As shown in [10], satisfiability for the Σ_2^0 class is in $\text{NTIME}((n \cdot k^m)^c)$, where n is the length of the input sentence, k and m are the number of existential and universal quantifiers respectively in the sentence, and c is a suitable constant. In recent years, there has been significant interest in the Σ_2^0 class from the program verification and program synthesis communities as well [25, 36, 62, 79]. Here, the Σ_2^0 class is also referred to as *effectively propositional logic*. For the Π_2^0 class on the other hand, the database community has shown a lot of active interest in this class in the context of data exchange, data integration and data interoperability [11, 26, 49, 53], and much more recently, in the context of query answering over RDF and OWL knowledge [43, 44].

In the remainder of this section, we describe our preservation properties, and our main results and techniques. All of these in the classical model theory setting are described in Section 1.1.1, and these in the finite model theory setting are described in Section 1.1.2. The latter section also contains our answer to the second part of Q2 raised above. The results that we present here contain, and generalize significantly, the results in [73, 75, 76].

1.1.1 Results in the classical model theory context

Our property of preservation under substructures modulo k -cruxes ($PSC(k)$), is a natural parameterized generalization of preservation under substructures, as can be seen from its definition (Definition 3.1.1): A sentence ϕ is $PSC(k)$ if every model \mathfrak{A} of ϕ contains a set C of at most k elements such that any substructure of \mathfrak{A} , *that contains C* , satisfies ϕ . It is evident that preservation under substructures is a special case of $PSC(k)$ when k equals 0. The property of preservation under k -ary covered extensions ($PCE(k)$) is defined as the dual of $PSC(k)$, whereby it generalizes the property of preservation under extensions (Definition 3.2.4). The generalized Los-Tarski theorem for sentences at level k ($\text{GLT}(k)$) gives syntactic characterizations of $PSC(k)$ and $PCE(k)$ as follows (Theorem 4.1.1): (i) an FO sentence is $PSC(k)$ if, and only if, it is equivalent to an $\exists^k \forall^*$ sentence, and (ii) an FO sentence is $PCE(k)$ if, and only if, it is equivalent to a $\forall^k \exists^*$ sentence. We call the former the *substructural version* of $\text{GLT}(k)$, and the latter the *extensional version* of $\text{GLT}(k)$. The Łoś-Tarski theorem for sentences is indeed a special case of $\text{GLT}(k)$ when k equals 0.

Towards extending $\text{GLT}(k)$ to the case of theories (sets of sentences), we first extend the notions of $PSC(k)$ and $PCE(k)$ to theories, and consider separately the substructural and extensional versions of $\text{GLT}(k)$. The extensional version of $\text{GLT}(k)$ lifts naturally: a theory is $PCE(k)$ if, and only if, it is equivalent to a theory of $\forall^k \exists^*$ sentences (Theorem 5.1.1(1)). The substructural version of $\text{GLT}(k)$ however does not lift to theories, as is witnessed by an intriguing counterexample that shows that there is a theory of $\exists \forall^*$ sentences, i.e. Σ_2^0 sentences with just one existential quantifier, that is not $PSC(k)$ for any k . Nevertheless, we show that $PSC(k)$ theories are always equivalent to Σ_2^0 theories, and as a (conditional) refinement of this result, we show that under a well-motivated model-theoretic hypothesis, $PSC(k)$ theories are equivalent to theories of $\exists^k \forall^*$ sentences (Theorems 5.2.1(2) and 5.2.3).

The above results give new semantic characterizations of the classes of Σ_2^0 and Π_2^0 sentences: the properties of “is $PSC(k)$ for some k ” and “is $PCE(k)$ for some k ” respectively characterize these sentences. The situation however becomes different when these characterizations are considered in the context of theories: Π_2^0 theories turn out to be more general than $PCE(k)$ theories for any k , and Σ_2^0 theories, indeed even $\exists \forall^*$ theories, turn out to be, as mentioned earlier, more general than $PSC(k)$ theories for any k . To get a characterization of Σ_2^0 and Π_2^0 theories by staying within the ambit of the flavour of our preservation properties, we introduce the properties of $PSC(\lambda)$ and $PCE(\lambda)$ as natural “infinitary” extensions of $PSC(k)$ and $PCE(k)$ respectively, in which the sizes of cruxes and arities of covers are now less than λ , for an infinite cardinal λ . Indeed, these extensions characterize Σ_2^0 and Π_2^0 theories (Theorems 5.1.1(2) and 5.2.1(1)), thereby giving new characterizations of the latter. We apply these characterizations to give new and simple proofs of well-known inexpressibility results in FO such as the inexpressibility of acyclicity, connectedness, bipartiteness, etc.

This completes the description of our results in the classical model theory context. We present various directions for future work, and sketch how natural generalizations of the properties of $PSC(k)$ and $PCE(k)$ can be used to get finer characterizations of Σ_n^0 and Π_n^0 sentences/theories for $n > 2$, analogous to the finer characterizations of Σ_2^0 and Π_2^0 sentences/theories by $PSC(k)$ and $PCE(k)$.

We conclude this subsection by briefly describing the techniques we use in proving our results described above. For $\text{GLT}(k)$, we give two proofs, one via a special class of structures called λ -saturated structures, and the other via ascending chains of structures (a similar proof works for the characterization of $PCE(k)$ and $PCE(\lambda)$ theories). To very quickly describe the for-

mer, we first show $\text{GLT}(k)$ over the class of λ -saturated structures, and then using the fact that any arbitrary structure has an “FO-similar” structure (i.e. a structure that satisfies the same FO sentences) that is λ -saturated, we “transfer” the truth of $\text{GLT}(k)$ over the class of λ -saturated structures, to that over the class of all structures. To show that $PSC(k)$ and $PSC(\lambda)$ theories are equivalent to Σ_2^0 theories, we use Keisler’s characterization of Σ_2^0 theories in terms of a preservation property defined in terms of 1-sandwiches, and show that any theory that is $PSC(k)$ or $PSC(\lambda)$ satisfies this preservation property. The proof of our result showing that under the well-motivated model-theoretic hypothesis alluded to earlier, a $PSC(k)$ theory is equivalent to an $\exists^k\forall^*$ theory, is the most involved of all our proofs. It introduces a novel technique of getting a syntactically defined FO theory equivalent to a given FO theory satisfying a semantic property, *by going outside of FO*. Specifically, for the case of $PSC(k)$ theories, under the aforementioned model-theoretic hypothesis, we first “go up” into an *infinitary logic* and show that a $PSC(k)$ theory can be characterized by syntactically defined sentences of this logic (Lemma 5.2.15). We then “come down” back to FO by providing a translation of the aforesaid infinitary sentences, to their equivalent FO theories, whenever these sentences are known to be equivalent to FO theories (Proposition 5.2.16). The FO theories are obtained from suitable *finite approximations* of the infinitary sentences, and turn out to be theories of $\exists^k\forall^*$ sentences. The “coming down” process can be seen as a “compilation” process (in the sense of compilers used in computer science) in which a “high level” description – via infinitary sentences that are known to be equivalent to FO theories – is translated into an equivalent “low level” description – via FO theories. We believe this technique of accessing the descriptive power of an infinitary logic followed by accessing the translation power of “compiler results” of the kind just mentioned, may have other applications as well.

1.1.2 Results in the finite model theory context

While the failure of the Łoś-Tarski theorem in the finite shows that universal sentences cannot capture in the finite, the property of preservation under substructures, we show a stronger result: that for any $k \geq 0$, the class of $\exists^k\forall^*$ also cannot capture in the finite, the property of preservation under substructures, and hence (not capture) $PSC(l)$ for any $l \geq 0$ (Proposition 8.1.1). This therefore shows the failure of $\text{GLT}(k)$ over all finite structures, for all $k \geq 0$. What happens to $\text{GLT}(k)$ over the well-behaved classes that have been identified by Atserias, Dawar and Grohe to admit the Łoś-Tarski theorem? It unfortunately turns out that none of the above classes,

in general, admits $\text{GLT}(k)$ for any $k \geq 2$. We show that the existence of induced paths of unbounded length in a class is, under reasonable assumptions on the class, the reason for the failure of $\text{GLT}(k)$ over the class (Theorem 8.2.2). Since these assumptions are satisfied by the aforesaid well-behaved classes and the latter allow unbounded induced path lengths in general, $\text{GLT}(k)$ fails over these classes in general.

To “recover” $\text{GLT}(k)$ in the face of the above failures, we define a new logic based combinatorial property of classes of finite structures, that we call the *\mathcal{L} -equivalent bounded substructure property*, denoted $\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$, where \mathcal{L} is either FO or an extension of FO, called *monadic second order logic* (MSO), \mathcal{S} is a class of finite structures, and k is a natural number (Definition 9.1). Intuitively, this property says that any structure \mathfrak{A} in \mathcal{S} contains a small substructure \mathfrak{B} that is in \mathcal{S} and that is “logically similar” to \mathfrak{A} . More precisely, \mathfrak{B} is “ (m, \mathcal{L}) -similar” to \mathfrak{A} , in that \mathfrak{B} and \mathfrak{A} agree on all \mathcal{L} sentences of quantifier rank m , where m is a given number. The bound on the size of \mathfrak{B} depends only on m (if \mathcal{S} and k are fixed). Further, such a small and (m, \mathcal{L}) -similar substructure can always be found “around” any given set of at most k elements of \mathfrak{A} .

We show that $\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ indeed entails $\text{GLT}(k)$. Interestingly, it also entails the homomorphism preservation theorem (HPT). (In fact, more general versions of $\text{GLT}(k)$ and HPT are entailed by $\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$; see Theorem 9.1.2 and Theorem 11.3.7.) Furthermore, if $\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ holds with “computable bounds”, i.e. if the bound on the size of the small substructure as referred to in the $\mathcal{L}\text{-EBSP}$ definition, is computable, then effective versions of the $\text{GLT}(k)$ and HPT are entailed by $\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$.

It turns out that a variety of classes of finite structures, that are of interest in computer science and finite model theory, satisfy $\mathcal{L}\text{-EBSP}(\cdot, k)$, and moreover, with computable bounds. The classes that we consider are broadly of two kinds: special kinds of labeled posets and special kinds of graphs. For the case of labeled posets, we show $\mathcal{L}\text{-EBSP}(\cdot, k)$ for the cases of words, trees (of various kinds such as unordered, ordered and ranked), and nested words over a finite alphabet, and regular subclasses of these (Theorem 10.2.2). While words and trees have had a long history of studies in the literature, nested words are much more recent [5], and have attracted a lot of attention as they admit a seamless generalization of the theory of regular languages and are also closely connected with visibly pushdown languages [4]. For the case of graphs, we show $\mathcal{L}\text{-EBSP}(\cdot, k)$ holds for a very general, and again very recently defined, class of graphs called *n -partite cographs*, and all hereditary subclasses of this class

(Theorem 10.3.1). This class of graphs, introduced in [31], jointly generalizes the classes of cographs, graph classes of bounded tree-depth and those of bounded shrub-depth. The latter graph classes have various interesting finiteness properties, and have become very prominent in the context of fixed parameter tractability of MSO model checking, and in the context of investigating when FO equals MSO in its expressive power [24, 28, 29, 52]. Being hereditary subclasses of the class of n -partite cographs, all these graph classes satisfy \mathcal{L} -EBSP(\cdot, k).

We go further to give many methods to construct new classes of structures satisfying \mathcal{L} -EBSP(\cdot, \cdot) (with computable bounds) from classes known to satisfy \mathcal{L} -EBSP(\cdot, \cdot) (with computable bounds). We show that \mathcal{L} -EBSP(\cdot, \cdot) is closed under taking subclasses that are hereditary or \mathcal{L} -definable, and is also closed under finite intersections and finite unions (Lemma 10.4.1). We show that \mathcal{L} -EBSP(\cdot, \cdot) remains preserved under various operations on structures, that have been well-studied in the literature: unary operations like complementation, transpose and the line graph operation, binary “sum-like” operations [57] such as disjoint union and join, and binary “product-like” operations that include various kinds of products like cartesian, tensor, lexicographic and strong products. All of these are examples of operations that can be implemented using what are called *quantifier-free translation schemes* [35, 57]. We show that FO-EBSP(\cdot, \cdot) is always closed under such operations, and MSO-EBSP(\cdot, k) is closed under such operations, provided that they are unary or sum-like. It follows that finite unions of classes obtained by finite compositions of the aforesaid operations also satisfies \mathcal{L} -EBSP(\cdot, \cdot). However, many interesting classes of structures can be obtained only by taking infinite unions of the kind just described, a notable example being the class of hamming graphs of the n -clique [42]. We show that if the aforementioned infinite unions are “regular”, in a sense we make precise, then \mathcal{L} -EBSP($\cdot, 0$) is preserved under these unions, under reasonable assumptions on the operations (Theorem 10.4.11). As applications of this result, we get that the class of hamming graphs of the n -clique satisfies FO-EBSP($\cdot, 0$), as does the class of p -dimensional grid posets, where p belongs to any MSO definable (using a linear order) class of natural numbers (like, even numbers).

The proofs of the above results rely on tree-representations of structures, and proceed by performing appropriate “prunings” of, and “graftings” within, these trees, in a manner that preserves the substructure and “ (m, \mathcal{L}) -similarity” relations between the structures represented by these trees. The process eventually yields small subtrees that represent bounded structures that are substructures of, and are (m, \mathcal{L}) -similar to, the original structures. Two key technical elements that are employed to perform the aforementioned prunings and graftings are the finiteness

of the index of the “ (m, \mathcal{L}) -similarity” relation (which is an equivalence relation) and the *type-transfer property* of the tree-representations. The latter means that the (m, \mathcal{L}) -similarity type of the structure represented by a tree t is *determined* by the multi-set of the (m, \mathcal{L}) -similarity types of the structures represented by the subtrees rooted at the children of the root of t , and further, determined only by a threshold number of appearances of each (m, \mathcal{L}) -similarity type in the multi-set, with the threshold depending solely on m . (Thus any change in the multi-set with respect to the (m, \mathcal{L}) -similarity types in it that appear less than threshold number of times, gets “transferred” to the (m, \mathcal{L}) -similarity type of the structure represented by the (changed) tree t .) These techniques have been incorporated into a single abstract result concerning tree representations, (Theorem 10.1.1), which we believe might be of independent interest.

Finally, we present three additional findings about the \mathcal{L} -EBSP(\cdot, k) property. We show that the \mathcal{L} -SAT problem (the problem of deciding if a given \mathcal{L} sentence is satisfiable) is decidable over any class satisfying \mathcal{L} -EBSP(\cdot, k) with computable bounds (Lemma 11.1.1). We next show that any class of structures that is well-quasi-ordered under the embedding relation satisfies \mathcal{L} -EBSP($\cdot, 0$) (Theorem 11.2.2). The notion of well-quasi-orders is very well-studied in the literature [39, 50, 59] and has great algorithmic implications. For instance, checking membership in any hereditary subclass of a well-quasi-ordered class can be done efficiently (i.e. in polynomial time). Our result above not only gives a technique to show the \mathcal{L} -EBSP($\cdot, 0$) property for a class (by showing the class to be well-quasi-ordered) but also, flipped around, gives a “logic-based” tool to show that a class of structures is not w.q.o. under embedding (by showing that the class does not satisfy \mathcal{L} -EBSP($\cdot, 0$)). Finally, we show that \mathcal{L} -EBSP(\cdot, k) can very well be seen to be a *finitary analogue* of the model-theoretic property that the classical downward Löwenheim-Skolem theorem (one the first results of classical model theory and a widely used tool in the subject, along with the compactness theorem) states of FO and arbitrary structures. This theorem says that an infinite structure \mathfrak{A} over a countable vocabulary always contains a countable substructure \mathfrak{B} that is “FO-similar” to \mathfrak{A} , in that \mathfrak{B} and \mathfrak{A} agree on all FO sentences. Further such a countable and FO-similar substructure can always be found “around” any given countable set of elements of \mathfrak{A} . The importance of the downward Löwenheim-Skolem theorem in classical model theory can be gauged from the fact that this theorem, along with the compactness theorem, characterizes FO [55]. It indeed is pleasantly surprising that while the downward Löwenheim-Skolem theorem is by itself meaningless over finite structures, a natural finitary analogue of the model-theoretic property that this theorem talks about, is satisfied by

a wide spectrum of classes of finite structures, that are of interest and importance in computer science and finite model theory.

This answers the second part of Q2 raised at the outset of Section 1.1.

We conclude this part of the thesis with several directions for future work, two of which we highlight here. The first asks for an investigation of a *structural characterization* of \mathcal{L} -EBSP(\cdot, k) motivated by the observation that any hereditary class of graphs satisfying \mathcal{L} -EBSP(\cdot, k) has bounded induced path lengths. The second of these is a conjecture. Though the failure of $\text{GLT}(k)$ over the class of all finite structures shows that $PSC(k)$ sentences are more expressive than $\exists^k \forall^*$ sentences over this class, it is still possible that all $PSC(k)$ sentences for all $k \geq 0$, taken together, are just as expressive as Σ_2^0 sentences, over all finite structures. We conjecture that this is indeed the case.

In summary, the properties and notions introduced in this thesis are interesting in both the classical and finite model theory contexts, and yield in both these contexts, a new, natural and parameterized generalization of the Łoś-Tarski preservation theorem.

1.2 Organization of the thesis

The thesis is in two parts: the first concerning classical model theory, and the second concerning finite model theory. We describe the organization within each of these parts below. For (the statements of) our results in the classical model theory part of the thesis (that are contained in Chapters 4, 5 and 6), we assume *arbitrary finite* vocabularies, unless explicitly stated otherwise (though the proofs of these results may resort to infinite vocabularies). In the finite model theory part of the thesis, we always consider *finite relational* vocabularies, unless we explicitly state otherwise.

Part I: Classical model theory

Chapter 2: We recall relevant notions and notation from classical model theory literature. We also recall the Łoś-Tarski theorem, and other results that we use in the subsequent chapters.

Chapter 3: We define the properties of $PSC(k)$ and $PCE(k)$, and formally show their duality.

Chapter 4: We characterize our properties and their variants for the case of sentences. We present $\text{GLT}(k)$, and provide two proofs of it using λ -saturated structures and ascending chains of structures (Section 4.1). We define the properties of $PSC(\lambda)$ and $PCE(\lambda)$ for infinite car-

dinals λ , provide characterizations of these properties, and present applications of these characterizations in getting new proofs of known inexpressibility results in FO (Section 4.2). We make further observations about our results so far, and prove an uncomputability result in connection with $PSC(k)$ (Section 4.3).

Chapter 5: We characterize all the properties introduced thus far, for the case of theories.

Chapter 6: We present directions for future work (in the classical model theory context).

Part II: Finite model theory

Chapter 7: We recall basic notions, notation, and results from the finite model theory literature.

Chapter 8: We discuss the need to investigate new classes of finite structures for $GLT(k)$.

Chapter 9: We define the property $\mathcal{L}\text{-EBSP}(\cdot, k)$, show that it entails $GLT(k)$ (Section 9.1) and show precisely, its connections with the downward Löwenheim-Skolem property (Section 9.2).

Chapter 10: We show that various classes of structures satisfy $\mathcal{L}\text{-EBSP}(\cdot, k)$. We first prove an abstract result concerning tree representations (Section 10.1), and then demonstrate its applications in showing the $\mathcal{L}\text{-EBSP}(\cdot, k)$ property for classes of words, trees (unordered, ordered, ranked) and nested words (Section 10.2), and the class of n -partite cographs and its various important subclasses (Section 10.3). We then give ways of constructing new classes satisfying $\mathcal{L}\text{-EBSP}(\cdot, \cdot)$, by presenting various closure properties of the latter (Section 10.4).

Chapter 11: We present additional studies on $\mathcal{L}\text{-EBSP}(\cdot, k)$. We show the decidability of the \mathcal{L} -satisfiability problem over classes satisfying $\mathcal{L}\text{-EBSP}(\cdot, k)$ with computable bounds (Section 11.1), and the connections of $\mathcal{L}\text{-EBSP}(\cdot, k)$ with well-quasi-orders (Section 11.2). We also show that $\mathcal{L}\text{-EBSP}(\cdot, k)$ entails a parameterized generalization of the homomorphism preservation theorem (Section 11.3).

Chapter 12: We present directions for future work (in the finite model theory context).

Chapter 13: We summarize the contributions of this thesis on both fronts, of classical model theory and finite model theory.

Part I

Classical Model Theory

Chapter 2

Background and preliminaries

In this part of the thesis, we shall be concerned with arbitrary structures (i.e. structures that are finite or infinite), and with the logic FO. The forthcoming sections of this chapter introduce the notation and terminology that we use throughout this part of the thesis. We also recall relevant results from the literature that we use in our proofs in the subsequent chapters. The classic references for all of the background that we set up in this chapter are [12, 40, 54].

We denote ordinals and cardinals using the letters λ, μ, κ or η . We let \mathbb{N} denote the set of natural numbers (zero included), and typically denote the elements of \mathbb{N} by the letters i, j etc. The cardinality of a set A is denoted as $|A|$; likewise the length of a tuple \bar{a} is denoted as $|\bar{a}|$. We denote $|\mathbb{N}|$ by either ω or \aleph_0 . We abbreviate in the standard way, some English phrases that commonly appear in mathematical literature. Specifically, ‘w.l.o.g’ stands for ‘without loss of generality’, ‘iff’ stands for ‘if and only if’, ‘w.r.t.’ stands for ‘with respect to’ and ‘resp.’ stands for ‘respectively’.

2.1 Syntax and semantics of FO

Syntax: A *vocabulary*, denoted by τ or σ , is a (possibly infinite) set of predicate, function and constant symbols. We denote variables by x, y, z , etc., possibly with numbers as subscripts. We denote a sequence of variables by $\bar{x}, \bar{y}, \bar{z}$, etc., again possibly with numbers as subscripts. We define below the notions of term, atomic formula and FO formula over τ .

1. A *term* over τ , or simply term if τ is clear from context, denoted using the letter ‘ t ’ typically along with numbers as subscripts, is defined inductively as follows:
 - (a) A constant (of τ) and a variable are terms each.

- (b) If t_1, \dots, t_n are terms over τ , then $f(t_1, \dots, t_n)$ is also a term over τ where f is an n -ary function symbol of τ .
2. An *atomic* formula over τ is one of the following:
- (a) The formula $t_1 = t_2$ where t_1 and t_2 are terms over τ and ‘=’ is a special predicate symbol not a part of τ which is interpreted always as the equality relation.
 - (b) The formula $R(t_1, \dots, t_n)$ where R is an n -ary relation symbol of τ , and t_1, \dots, t_n are terms over τ .
3. An *FO formula* over τ , also called an $\text{FO}(\tau)$ formula, or simply formula if τ is clear from context, is defined inductively as follows:
- (a) An atomic formula over τ is an $\text{FO}(\tau)$ formula.
 - (b) If φ_1 and φ_2 are $\text{FO}(\tau)$ formulae, then each of $\varphi_1 \wedge \varphi_2$, $\varphi_1 \vee \varphi_2$ and $\neg \varphi_1$ are also $\text{FO}(\tau)$ formulae. Here, the symbols \wedge, \vee and \neg denote the usual boolean connectives ‘and’, ‘or’ and ‘not’ respectively.
 - (c) If φ_1 is an $\text{FO}(\tau)$ formula, then $\exists x \varphi_1$ and $\forall x \varphi_1$ are also $\text{FO}(\tau)$ formulae. Here the symbols \exists and \forall denote respectively, the existential and universal quantifiers.

In addition to the letter φ , we use other Greek letters like $\phi, \psi, \chi, \xi, \gamma, \alpha$ and β to denote formulae. A formula without any quantifiers is called *quantifier-free*. We abbreviate a block of quantifiers of the form $Qx_1 \dots Qx_k$ by $Q^k \bar{x}$ or $Q\bar{x}$ (depending on what is better suited for the context), where $Q \in \{\forall, \exists\}$ and $k \in \mathbb{N}$. By Q^* , we mean a block of k Q quantifiers, for some $k \in \mathbb{N}$.

We now define the notion of *free variables* of a term or a formula. The term x , where x is a variable, has only one free variable, which is x itself. A term that is a constant has no free variables. The set of free variables of the term $f(t_1, \dots, t_n)$ is the union of the sets of free variables of t_1, \dots, t_n . The latter is also the set of free variables of the atomic formula $R(t_1, \dots, t_n)$. Any free variable of the atomic formula $t_1 = t_2$ is a free variable of either t_1 or t_2 . The set of free variables of the formula $\varphi_1 \wedge \varphi_2$, and of the formula $\varphi_1 \vee \varphi_2$, is the union of the sets of free variables of φ_1 and φ_2 . Negation preserves the free variables of a formula. Finally, the free variables of $\exists x \varphi$ and $\forall x \varphi$ are the free variables of φ except for x . We let $t(\bar{x})$, resp. $\varphi(\bar{x})$, denote a term t , resp. formula φ , whose free variables are *among* \bar{x} . A formula with no free variables is called a *sentence*.

Semantics: Let $\tau = \tau_C \sqcup \tau_R \sqcup \tau_F$ where τ_C, τ_R and τ_F are respectively the set of constant, relation and function symbols of τ . A τ -*structure* $\mathfrak{A} = (\mathcal{U}_{\mathfrak{A}}, (c^{\mathfrak{A}})_{c \in \tau_C}, (R^{\mathfrak{A}})_{R \in \tau_R}, (f^{\mathfrak{A}})_{f \in \tau_F})$

consists of a set $U_{\mathfrak{A}}$ called the *universe* or the *domain* of \mathfrak{A} , along with interpretations $c^{\mathfrak{A}}$, $R^{\mathfrak{A}}$ and $f^{\mathfrak{A}}$ of each of the symbols c , R and f of τ_C , τ_R and τ_F respectively, such that

- the constant symbol c is interpreted as an element $c^{\mathfrak{A}} \in U_{\mathfrak{A}}$
- the n -ary relation symbol R is interpreted as a set $R^{\mathfrak{A}}$ of n -tuples of \mathfrak{A} , i.e. $R^{\mathfrak{A}} \subseteq (U_{\mathfrak{A}})^n$
- the n -ary function symbol f is interpreted as a function $f^{\mathfrak{A}} : (U_{\mathfrak{A}})^n \rightarrow U_{\mathfrak{A}}$

When τ is clear from context, we refer to a τ -structure as simply a structure. We denote structures by \mathfrak{A} , \mathfrak{B} etc.

Towards the semantics of FO, we first define, for a given structure \mathfrak{A} and a term $t(x_1, \dots, x_n)$, the *value* in \mathfrak{A} , of $t(x_1, \dots, x_n)$ for a given assignment a_1, \dots, a_n of elements of \mathfrak{A} , to the variables x_1, \dots, x_n . We denote this value as $t^{\mathfrak{A}}(\bar{a})$ where $\bar{a} = (a_1, \dots, a_n)$.

- If t is a constant symbol, then $t^{\mathfrak{A}}(\bar{a}) = c^{\mathfrak{A}}$.
- If t is the variable x_i , then $t^{\mathfrak{A}}(\bar{a}) = a_i$.
- If $t = f(t_1, \dots, t_n)$, then $t^{\mathfrak{A}}(\bar{a}) = f^{\mathfrak{A}}(t_1^{\mathfrak{A}}(\bar{a}), \dots, t_n^{\mathfrak{A}}(\bar{a}))$.

We now define, for a given structure \mathfrak{A} and formula $\varphi(\bar{x})$, the notion of the *truth* of $\varphi(\bar{x})$ in \mathfrak{A} given an assignment \bar{a} of elements of \mathfrak{A} , to the variables \bar{x} . We denote by $(\mathfrak{A}, \bar{a}) \models \varphi(\bar{x})$, that $\varphi(\bar{x})$ is true in \mathfrak{A} for the assignment \bar{a} to \bar{x} . We then call (\mathfrak{A}, \bar{a}) a *model* of $\varphi(\bar{x})$.

- If $\varphi(\bar{x})$ is the formula $t_1 = t_2$, then $(\mathfrak{A}, \bar{a}) \models \varphi(\bar{x})$ iff $t_1^{\mathfrak{A}}(\bar{a}) = t_2^{\mathfrak{A}}(\bar{a})$.
- If $\varphi(\bar{x})$ is the formula $R(t_1, \dots, t_n)$, then $(\mathfrak{A}, \bar{a}) \models \varphi(\bar{x})$ iff $(t_1^{\mathfrak{A}}(\bar{a}), \dots, t_n^{\mathfrak{A}}(\bar{a})) \in R^{\mathfrak{A}}$.
- If $\varphi(\bar{x})$ is the formula $\varphi_1(\bar{x}) \wedge \varphi_2(\bar{x})$, then $(\mathfrak{A}, \bar{a}) \models \varphi(\bar{x})$ iff $(\mathfrak{A}, \bar{a}) \models \varphi_1(\bar{x})$ and $(\mathfrak{A}, \bar{a}) \models \varphi_2(\bar{x})$.
- If $\varphi(\bar{x})$ is the formula $\varphi_1(\bar{x}) \vee \varphi_2(\bar{x})$, then $(\mathfrak{A}, \bar{a}) \models \varphi(\bar{x})$ iff $(\mathfrak{A}, \bar{a}) \models \varphi_1(\bar{x})$ or $(\mathfrak{A}, \bar{a}) \models \varphi_2(\bar{x})$.
- If $\varphi(\bar{x})$ is the formula $\neg \varphi_1(\bar{x})$, then $(\mathfrak{A}, \bar{a}) \models \varphi(\bar{x})$ iff it is not the case that $(\mathfrak{A}, \bar{a}) \models \varphi_1(\bar{x})$.
- If $\varphi(\bar{x})$ is the formula $\exists y \varphi_1(\bar{x}, y)$, then $(\mathfrak{A}, \bar{a}) \models \varphi(\bar{x})$ iff there exists an element $b \in U_{\mathfrak{A}}$ such that $(\mathfrak{A}, \bar{a}, b) \models \varphi_1(\bar{x}, y)$.
- If $\varphi(\bar{x})$ is the formula $\forall y \varphi_1(\bar{x}, y)$, then $(\mathfrak{A}, \bar{a}) \models \varphi(\bar{x})$ iff for all elements $b \in U_{\mathfrak{A}}$, $(\mathfrak{A}, \bar{a}, b) \models \varphi_1(\bar{x}, y)$.

If φ is a sentence, we denote the truth of φ in \mathfrak{A} simply as $\mathfrak{A} \models \varphi$, and call \mathfrak{A} a model of φ .

Given an $\text{FO}(\tau)$ formula $\varphi(x_1, \dots, x_n)$ and distinct constants c_1, \dots, c_n not appearing in τ , let φ' be the FO sentence over the vocabulary $\tau \cup \{c_1, \dots, c_n\}$, obtained by substituting c_i for the free occurrences of x_i in $\varphi(x_1, \dots, x_n)$ for each $i \in \{1, \dots, n\}$. The following lemma connects

the notions of truth of $\varphi(x_1, \dots, x_n)$ in a model and the truth of φ' in a model.

Lemma 2.1.1. $(\mathfrak{A}, a_1, \dots, a_n) \models \varphi(x_1, \dots, x_n)$ iff $(\mathfrak{A}, a_1, \dots, a_n) \models \varphi'$.

Note the distinction between the two occurrences of “ $(\mathfrak{A}, a_1, \dots, a_n)$ ” in the lemma above. The occurrence on the left denotes that a_1, \dots, a_n is an assignment to x_1, \dots, x_n in the τ -structure \mathfrak{A} , whereas the occurrence on the right denotes a τ_n -structure.

Extending syntax and semantics to theories: In classical model theory, one frequently talks about FO theories. We define the syntax and semantics of these now. A *theory*, resp. $FO(\tau)$ *theory*, is simply a set of sentences, resp. a set of $FO(\tau)$ sentences. An $FO(\tau)$ theory is also referred to as a *theory over τ* . We typically denote theories using capital letters like T, V, Y, Z , possibly with numbers as subscripts. A theory, resp. $FO(\tau)$ theory, *whose free variables are among \bar{x}* , is a set of formulae, resp. $FO(\tau)$ formulae, all of whose free variables are among \bar{x} . We denote such theories as $T(\bar{x}), V(\bar{x})$ etc. Given a theory $T(\bar{x})$, a structure \mathfrak{A} , and a tuple \bar{a} from \mathfrak{A} such that $|\bar{a}| = |\bar{x}|$, we denote by $(\mathfrak{A}, \bar{a}) \models T(\bar{x})$ that $(\mathfrak{A}, \bar{a}) \models \varphi(\bar{x})$ for each formula $\varphi(\bar{x}) \in T(\bar{x})$. In such a case, we say $T(\bar{x})$ is *true* in \mathfrak{A} for the assignment \bar{a} to \bar{x} , and that (\mathfrak{A}, \bar{a}) is a *model* of $T(\bar{x})$. If T has no free variables, then we denote the truth of T in \mathfrak{A} as $\mathfrak{A} \models T$, and say that \mathfrak{A} is a model of T . Observe that if T is empty, then trivially, every structure \mathfrak{A} is a model of T .

Given an $FO(\tau)$ theory $T(x_1, \dots, x_n)$ and distinct constants c_1, \dots, c_n not appearing in τ , let T' be the FO theory without free variables over the vocabulary $\tau \cup \{c_1, \dots, c_n\}$, obtained by substituting c_i for the free occurrences of x_i in $T(x_1, \dots, x_n)$ for each $i \in \{1, \dots, n\}$. Analogous to Lemma 2.1.1, we have the following lemma.

Lemma 2.1.2. $(\mathfrak{A}, a_1, \dots, a_n) \models T(x_1, \dots, x_n)$ iff $(\mathfrak{A}, a_1, \dots, a_n) \models T'$.

Consistency, validity, entailment and equivalence: Let $T(\bar{x})$ be a given theory and $\varphi(\bar{x})$ be a given formula. We say $T(\bar{x})$ is *consistent* or *satisfiable*¹ if it has a model, i.e. if there exists a structure \mathfrak{A} and tuple \bar{a} of \mathfrak{A} such that $|\bar{a}| = |\bar{x}|$ and $(\mathfrak{A}, \bar{a}) \models T(\bar{x})$. If $T(\bar{x})$ is not consistent, then we say it is *inconsistent* or *unsatisfiable*. We say $T(\bar{x})$ is *valid* if $(\mathfrak{A}, \bar{a}) \models T(\bar{x})$ for every structure \mathfrak{A} and every tuple \bar{a} of \mathfrak{A} such that $|\bar{a}| = |\bar{x}|$. The notions above have natural adaptations to formulae. We say $\varphi(\bar{x})$ is satisfiable, unsatisfiable, or valid if $\{\varphi(\bar{x})\}$ is

¹In the literature, consistency has a proof-theoretic definition. However Gödel’s completeness theorem shows that for FO, consistency is the same as satisfiability, the latter meaning the existence of a model. Hence, we do not make a distinction between consistency and satisfiability in this thesis.

satisfiable, unsatisfiable, or valid, respectively. It is easy to see that $\varphi(\bar{x})$ is valid iff $\neg\varphi(\bar{x})$ is unsatisfiable, and that $\varphi(\bar{x})$ and $\neg\varphi(\bar{x})$ can both be satisfiable.

We say $T(\bar{x})$ *entails* $\varphi(\bar{x})$, denoted $T(\bar{x}) \vdash \varphi(\bar{x})$, if every model (\mathfrak{A}, \bar{a}) of $T(\bar{x})$ is also a model of $\varphi(\bar{x})$. For a formula $\psi(\bar{x})$, we denote by $\psi(\bar{x}) \rightarrow \varphi(\bar{x})$, that the theory $\{\psi(\bar{x})\}$ entails $\varphi(\bar{x})$. It is easy to verify that $\psi(\bar{x}) \rightarrow \varphi(\bar{x})$ iff $\neg\psi(\bar{x}) \vee \varphi(\bar{x})$ is valid. Given a theory $Y(\bar{x})$, we say $T(\bar{x})$ is *equivalent to* $Y(\bar{x})$ if $T(\bar{x})$ entails every formula of $Y(\bar{x})$, and vice-versa. We denote by $\psi(\bar{x}) \leftrightarrow \varphi(\bar{x})$ that $\{\psi(\bar{x})\}$ is equivalent to $\{\varphi(\bar{x})\}$.

We now adapt all the notions above to versions of these *modulo theories*. Given a consistent theory V , we say $T(\bar{x})$ is *consistent or satisfiable modulo* V if $(V \cup T(\bar{x}))$ is consistent, and say $T(\bar{x})$ is *inconsistent or unsatisfiable modulo* V if $(V \cup T(\bar{x}))$ is inconsistent. We say $T(\bar{x})$ *entails* $\varphi(\bar{x})$ *modulo* V if $(V \cup T(\bar{x})) \vdash \varphi(\bar{x})$, and say $T(\bar{x})$ and $Y(\bar{x})$ are *equivalent modulo* V if $(V \cup T(\bar{x}))$ is equivalent to $(V \cup Y(\bar{x}))$. This last notion is particularly relevant for this part of the thesis, and stated in other words, it says that for every model \mathfrak{A} of V , and for every tuple \bar{a} from \mathfrak{A} such that $|\bar{a}| = |\bar{x}|$, it is the case that $(\mathfrak{A}, \bar{a}) \models T(\bar{x})$ iff $(\mathfrak{A}, \bar{a}) \models Y(\bar{x})$. One can define all the notions just mentioned, for formulae, analogously as in the previous paragraphs.

2.2 Σ_n^0 and Π_n^0 formulae

An FO formula in which all quantifiers appear first (from left to right) followed by a quantifier-free formula, is said to be in *prenex normal form*. For such a formula, the sequence of quantifiers is called the *quantifier prefix*, and the quantifier-free part is called the *matrix* of the formula. For every non-zero $n \in \mathbb{N}$, we denote by Σ_n^0 , resp. Π_n^0 , the class of all FO formulae in prenex normal form, whose quantifier prefix begins with \exists , resp. \forall , and consists of $n - 1$ alternations of quantifiers. We call Σ_1^0 formulae *existential* and Π_1^0 formulae *universal*. We call Σ_2^0 formulae having k existential quantifiers $\exists^k \forall^*$ *formulae*, and Π_2^0 formulae having k universal quantifiers $\forall^k \exists^*$ *formulae*. The aforementioned notions for formulae, have natural liftings to theories: A Σ_n^0 theory, resp. Π_n^0 theory, is a theory all of whose formulae are Σ_n^0 , resp. Π_n^0 ; an existential theory, resp. universal theory, is a Σ_1^0 theory, resp. Π_1^0 theory; an $\exists^k \forall^*$ theory, resp. $\forall^k \exists^*$ theory, is a theory all of whose formulae are $\exists^k \forall^*$, resp. $\forall^k \exists^*$. We now have the following lemma.

Lemma 2.2.1. *Every FO formula is equivalent to an FO formula in prenex normal form. By extension, every FO theory is equivalent to a theory of FO formulae, all of which are in prenex normal form.*

2.3 Notions concerning structures

The *size* (or *power*) of a structure \mathfrak{A} is the cardinality of $U_{\mathfrak{A}}$. A structure is called *finite* if its size is finite, else it is called *infinite*. A *substructure of \mathfrak{A} induced by a subset B of $U_{\mathfrak{A}}$* is a structure \mathfrak{B} such that (i) $U_{\mathfrak{B}} = \{t^{\mathfrak{A}}(\bar{a}) \mid t(x_1, \dots, x_n) \text{ is a term over } \tau, \bar{a} \text{ is an } n\text{-tuple from } B\}$ (ii) $c^{\mathfrak{B}} = c^{\mathfrak{A}}$ for each constant symbol $c \in \tau$, (iii) $R^{\mathfrak{B}} = R^{\mathfrak{A}} \cap (U_{\mathfrak{B}})^n$ for each n -ary relation symbol $R \in \tau$, and (iv) $f^{\mathfrak{B}}$ is the restriction of $f^{\mathfrak{A}}$ to $U_{\mathfrak{B}}$, for each n -ary function symbol $f \in \tau$. A *substructure \mathfrak{B} of \mathfrak{A}* , denoted $\mathfrak{B} \subseteq \mathfrak{A}$, is a substructure of \mathfrak{A} induced by some subset of $U_{\mathfrak{A}}$. If $\mathfrak{B} \subseteq \mathfrak{A}$, we say \mathfrak{A} is an *extension* of \mathfrak{B} . It is easy to see that if $\mathfrak{B} \subseteq \mathfrak{A}$, then for all quantifier-free formulae $\varphi(\bar{x})$ and all n -tuples \bar{a} from \mathfrak{B} where $n = |\bar{x}|$, we have that $(\mathfrak{B}, \bar{a}) \models \varphi(\bar{x})$ iff $(\mathfrak{A}, \bar{a}) \models \varphi(\bar{x})$. If (\mathfrak{B}, \bar{a}) and (\mathfrak{A}, \bar{a}) agree on *all* FO formulae $\varphi(\bar{x})$ (instead of only quantifier-free formulae) for all n -tuples \bar{a} from \mathfrak{B} where $n = |\bar{x}|$, then we say \mathfrak{B} is an *elementary substructure* of \mathfrak{A} , or \mathfrak{A} is an *elementary extension* of \mathfrak{B} , and denote it as $\mathfrak{B} \preceq \mathfrak{A}$. A notion related to the notion of elementary substructure is that of elementary equivalence: We say two structures \mathfrak{A} and \mathfrak{C} are *elementarily equivalent*, denoted $\mathfrak{A} \equiv \mathfrak{C}$, if they agree on all FO sentences.

Given τ -structures \mathfrak{A} and \mathfrak{B} , an *isomorphism* from \mathfrak{A} to \mathfrak{B} , denoted $h : \mathfrak{A} \rightarrow \mathfrak{B}$, is a bijection $h : U_{\mathfrak{A}} \rightarrow U_{\mathfrak{B}}$ such that (i) $c^{\mathfrak{B}} = h(c^{\mathfrak{A}})$ for every constant symbol $c \in \tau$ (ii) $(a_1, \dots, a_n) \in R^{\mathfrak{A}}$ iff $(h(a_1), \dots, h(a_n)) \in R^{\mathfrak{B}}$ for every n -ary relation symbol $R \in \tau$, and (iii) $f^{\mathfrak{A}}(a_1, \dots, a_n) = a$ iff $f^{\mathfrak{B}}(h(a_1), \dots, h(a_n)) = h(a)$ for every n -ary function symbol $f \in \tau$. If an isomorphism from \mathfrak{A} to \mathfrak{B} exists, then so does an isomorphism from \mathfrak{B} to \mathfrak{A} (namely, the inverse of the former isomorphism), and we say \mathfrak{A} and \mathfrak{B} are *isomorphic*, and denote it as $\mathfrak{A} \cong \mathfrak{B}$. We say \mathfrak{A} is *(isomorphically) embeddable* in \mathfrak{B} , or simply *embeddable* in \mathfrak{B} , denoted as $\mathfrak{A} \hookrightarrow \mathfrak{B}$, if there exists a substructure \mathfrak{C} of \mathfrak{B} such that there is an isomorphism $h : \mathfrak{A} \rightarrow \mathfrak{C}$. In such a case, we say h is an *(isomorphic) embedding*, or simply an *embedding*, of \mathfrak{A} in \mathfrak{B} . We say \mathfrak{A} is *elementarily embeddable* in \mathfrak{B} if there exists an elementary substructure \mathfrak{C} of \mathfrak{B} such that there is an isomorphism $h : \mathfrak{A} \rightarrow \mathfrak{C}$. In such a case, we say h is an *elementary embedding* of \mathfrak{A} in \mathfrak{B} . Given vocabularies τ, τ' , we say τ' is an *expansion* of τ if $\tau \subseteq \tau'$. Given a τ -structure \mathfrak{A} and a τ' -structure \mathfrak{A}' , we say \mathfrak{A}' is a τ' -*expansion* of \mathfrak{A} , or simply an *expansion* of \mathfrak{A} (if τ' is clear from context), if the universe of \mathfrak{A}' and the interpretations in \mathfrak{A}' , of the constant, predicate and function symbols of τ are exactly the same as those in \mathfrak{A} respectively. In such a case, we also say that \mathfrak{A} is a τ -*reduct* of \mathfrak{A}' . In this thesis, we will mostly consider expansions τ' of τ in which all the symbols of $\tau' \setminus \tau$ are constants. Given a cardinal λ , we denote by τ_{λ} , a fixed expansion of

τ such that $\tau_\lambda \setminus \tau$ consists only of the constants c_1, \dots, c_λ that are distinct and do not appear in τ , and say that τ_λ is an expansion of τ with constants c_1, \dots, c_λ . Given a τ -structure \mathfrak{A} and a λ -tuple (i.e. a tuple of length λ) $\bar{a} = (a_1, \dots, a_\lambda)$ of elements of \mathfrak{A} , we denote by (\mathfrak{A}, \bar{a}) the τ_λ -structure whose τ -reduct is \mathfrak{A} , and in which c_i is interpreted as a_i , for each $i < \lambda$. Given a τ -structure \mathfrak{A} and a subset X of $U_{\mathfrak{A}}$, we denote by τ_X , the expansion of τ with $|X|$ many fresh and distinct constants, one constant per element of X . We denote by $(\mathfrak{A}, (a)_{a \in X})$ the τ_X -expansion of \mathfrak{A} in which the constant in $\tau_X \setminus \tau$ corresponding to an element a of X , is interpreted as a itself. Given a τ -structure \mathfrak{B} such that $\mathfrak{B} \subseteq \mathfrak{A}$, if $X = U_{\mathfrak{B}}$, then we denote τ_X as $\tau_{\mathfrak{B}}$, and the structure $(\mathfrak{A}, (a)_{a \in X})$ as $\mathfrak{A}_{\mathfrak{B}}$. The *diagram* of \mathfrak{A} , denoted $\text{Diag}(\mathfrak{A})$, is the set of all quantifier-free $\text{FO}(\tau_{\mathfrak{A}})$ sentences that are true in $\mathfrak{A}_{\mathfrak{A}}$. The *elementary diagram* of \mathfrak{A} , denoted $\text{El-diag}(\mathfrak{A})$, is the set of all $\text{FO}(\tau_{\mathfrak{A}})$ sentences that are true in $\mathfrak{A}_{\mathfrak{A}}$. The following lemma connects the notions described in the previous and current paragraphs.

Lemma 2.3.1. *Let \mathfrak{A} and \mathfrak{B} be given τ -structures. Then the following are true.*

1. $\mathfrak{A} \cong \mathfrak{B}$ implies $\mathfrak{A} \equiv \mathfrak{B}$.
2. $\mathfrak{A} \preceq \mathfrak{B}$ implies $\mathfrak{A} \equiv \mathfrak{B}$.
3. $\mathfrak{A} \preceq \mathfrak{B}$ iff $\mathfrak{A}_{\mathfrak{A}} \equiv \mathfrak{B}_{\mathfrak{A}}$.
4. \mathfrak{A} is embeddable in \mathfrak{B} iff for some $\tau_{\mathfrak{A}}$ -expansion \mathfrak{B}' of \mathfrak{B} , it is the case that $\mathfrak{B}' \models \text{Diag}(\mathfrak{A})$.
5. \mathfrak{A} is elementarily embeddable in \mathfrak{B} iff for some $\tau_{\mathfrak{A}}$ -expansion \mathfrak{B}' of \mathfrak{B} , it is the case that $\mathfrak{B}' \models \text{El-diag}(\mathfrak{A})$.
6. If \mathfrak{A} is finite and $\mathfrak{A} \equiv \mathfrak{B}$, then $\mathfrak{A} \cong \mathfrak{B}$.

A class of structures is said to be *elementary* if it is the class of models of an FO theory. It is easy to see that an elementary class of structures is closed under elementary equivalence, and hence under isomorphisms.

We conclude this section by recalling some important results from the literature [12]. The first two of these below are arguably the most important theorems² of classical model theory (see Theorem 1.3.22 in [12] and Corollary 3.1.4 in [40]).

Theorem 2.3.2 (Compactness theorem, Gödel 1930, Mal'tsev 1936). *A theory $T(\bar{x})$ is consistent iff every finite subset of it is consistent.*

²As an aside, by a celebrated result of Lindström [55], FO is the only logic having certain well-defined and reasonable closure properties, that satisfies Theorem 2.3.2 and Theorem 2.3.3.

Theorem 2.3.3 (Downward Löwenheim-Skolem theorem, Löwenheim 1915, Skolem 1920s, Mal'tsev 1936). *Let \mathfrak{A} be a structure over a countable vocabulary and W be a set of at most λ elements of \mathfrak{A} , where λ is an infinite cardinal. Then there exists an elementary substructure \mathfrak{B} of \mathfrak{A} , that contains W and that has size at most λ .*

An easy but important corollary of the compactness theorem is the following.

Lemma 2.3.4 (Corollary 5.4.2, ref. [40]). *Let \mathfrak{A} and \mathfrak{B} be structures such that every existential sentence that is true in \mathfrak{B} is true in \mathfrak{A} . Then \mathfrak{B} is embeddable in an elementary extension of \mathfrak{A} .*

To state the final result that we recall here from literature, we need some terminology. Given a cardinal λ , an *ascending chain*, or simply *chain*, $(\mathfrak{A}_\eta)_{\eta < \lambda}$ of structures is a sequence $\mathfrak{A}_0, \mathfrak{A}_1, \dots$ of structures such that $\mathfrak{A}_0 \subseteq \mathfrak{A}_1 \subseteq \dots$. The *union* of the chain $(\mathfrak{A}_\eta)_{\eta < \lambda}$ is a structure \mathfrak{A} defined as follows: (i) $U_{\mathfrak{A}} = \bigcup_{\eta < \lambda} U_{\mathfrak{A}_\eta}$ (ii) $c^{\mathfrak{A}} = c^{\mathfrak{A}_\eta}$ for every constant symbol $c \in \tau$ and every $\eta < \lambda$ (observe that $c^{\mathfrak{A}}$ is well-defined) (iii) $R^{\mathfrak{A}} = \bigcup_{\eta < \lambda} R^{\mathfrak{A}_\eta}$ for every relation symbol $R \in \tau$ (iv) $f^{\mathfrak{A}} = \bigcup_{\eta < \lambda} f^{\mathfrak{A}_\eta}$ for every function symbol $f \in \tau$ (here, in taking the union of functions, we view an n -ary function as its corresponding $(n+1)$ -ary relation). It is clear that \mathfrak{A} is well-defined. We denote \mathfrak{A} as $\bigcup_{\eta < \lambda} \mathfrak{A}_\eta$. A chain $(\mathfrak{A}_\eta)_{\eta < \lambda}$ with the property that $\mathfrak{A}_0 \preceq \mathfrak{A}_1 \preceq \dots$ is said to be an *elementary chain*. We now have the following result (Theorem 3.1.9 of [12]).

Theorem 2.3.5 (Elementary chain theorem, Tarski-Vaught). *Let $(\mathfrak{A}_\eta)_{\eta < \lambda}$ be an elementary chain of structures. Then $\bigcup_{\eta < \lambda} \mathfrak{A}_\eta$ is an elementary extension of \mathfrak{A}_η for each $\eta < \lambda$.*

2.4 Types and μ -saturation

Given a vocabulary τ , a set $\Gamma(x_1, \dots, x_k)$ of $\text{FO}(\tau)$ formulae, all of whose free variables are among x_1, \dots, x_k , is said to be an *FO-type* of τ , or simply a *type* of τ , if it is *maximally consistent*, i.e. if it is consistent and for any $\text{FO}(\tau)$ formula $\varphi(x_1, \dots, x_k)$, exactly one of $\varphi(x_1, \dots, x_k)$ and $\neg\varphi(x_1, \dots, x_k)$ belongs to $\Gamma(x_1, \dots, x_k)$. Given a τ -structure \mathfrak{A} and a k -tuple \bar{a} of \mathfrak{A} , we let $\text{tp}_{\mathfrak{A}, \bar{a}}(x_1, \dots, x_k)$ denote the *type of \bar{a} in \mathfrak{A}* , i.e. the set of all $\text{FO}(\tau)$ formulae $\varphi(x_1, \dots, x_k)$ such that $(\mathfrak{A}, \bar{a}) \models \varphi(x_1, \dots, x_k)$. It is clear that $\text{tp}_{\mathfrak{A}, \bar{a}}(x_1, \dots, x_k)$ is a type of τ . The converse is clear too: if $\Gamma(x_1, \dots, x_k)$ is a type of τ , then for some τ -structure \mathfrak{A} and some k -tuple \bar{a} of \mathfrak{A} , it is the case that $\Gamma(x_1, \dots, x_k)$ is the type of \bar{a} in \mathfrak{A} . In such a case, we say that \mathfrak{A} *realizes* $\Gamma(x_1, \dots, x_k)$, and that \bar{a} *satisfies*, or *realizes*, $\Gamma(x_1, \dots, x_k)$ in \mathfrak{A} . It is easy to see for given structures \mathfrak{A} and \mathfrak{B} , and given k -tuples \bar{a} and \bar{b} from \mathfrak{A} and \mathfrak{B} resp., that $\text{tp}_{\mathfrak{A}, \bar{a}}(x_1, \dots, x_k)$

$= \text{tp}_{\mathfrak{A}, \bar{a}}(x_1, \dots, x_k)$ iff $(\mathfrak{A}, \bar{a}) \equiv (\mathfrak{B}, \bar{b})$. By $\text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, \dots, x_k)$, we denote the Π_1^0 -type of \bar{a} in \mathfrak{A} , i.e. the subset of $\text{tp}_{\mathfrak{A}, \bar{a}}(x_1, \dots, x_k)$ that consists of all Π_1^0 formulae of the latter. We denote by $\text{Th}(\mathfrak{A})$, the *theory of \mathfrak{A}* , i.e. the set of all $\text{FO}(\tau)$ sentences that are true in \mathfrak{A} .

We now recall the important notion of μ -saturated structures from the literature.

Definition 2.4.1 (Chp. 5, ref. [12]). Let μ be a cardinal. A τ -structure \mathfrak{A} is said to be μ -saturated if for every subset X of $\text{U}_{\mathfrak{A}}$, of cardinality less than μ , if \mathfrak{A}' is the τ_X -expansion $(\mathfrak{A}, (a)_{a \in X})$ of \mathfrak{A} , then \mathfrak{A}' realizes every type $\Gamma(x)$ of the vocabulary τ_X , that is consistent modulo $\text{Th}(\mathfrak{A}')$.

Following are some results in connection with μ -saturated structures, that we crucially use in many proofs in the forthcoming chapters.

Proposition 2.4.2. *The following are true for any vocabulary τ and any τ -structure \mathfrak{A} .*

1. [Proposition 5.1.1, ref. [12]] \mathfrak{A} is μ -saturated if and only if for every ordinal $\eta < \mu$ and every η -tuple \bar{a} of \mathfrak{A} , the expansion (\mathfrak{A}, \bar{a}) is μ -saturated.
2. [Proposition 5.1.2, ref. [12]] \mathfrak{A} is finite if and only if \mathfrak{A} is μ -saturated for all cardinals μ .
3. [Lemma 5.1.4, ref. [12]] There exists a μ -saturated elementary extension of \mathfrak{A} , for some cardinal $\mu \geq |\tau|$.
4. [Lemma 5.1.10, ref. [12]] If \mathfrak{A} is μ -saturated, $\mathfrak{A} \equiv \mathfrak{B}$ and \bar{b} is an η -tuple of \mathfrak{B} where $\eta < \mu$, then there exists an η -tuple \bar{a} of \mathfrak{A} such that $(\mathfrak{A}, \bar{a}) \equiv (\mathfrak{B}, \bar{b})$.
5. [Lemma 5.2.1, ref. [12]] Suppose every existential sentence that holds in \mathfrak{A} also holds in \mathfrak{B} , where \mathfrak{B} is μ -saturated for $\mu \geq |\mathfrak{A}|$. Then \mathfrak{A} is embeddable in \mathfrak{B} .

2.5 Two classical preservation properties

We first recall the classical dual notions of preservation under substructures and preservation under extensions. We fix a finite vocabulary τ in our discussion below.

Definition 2.5.1. Let \mathcal{S} be a class of structures.

1. A subclass \mathcal{U} of \mathcal{S} is said to be *preserved under substructures over \mathcal{S}* , abbreviated as \mathcal{U} is *PS over \mathcal{S}* , if for each structure $\mathfrak{A} \in \mathcal{U}$, if $\mathfrak{B} \subseteq \mathfrak{A}$ and $\mathfrak{B} \in \mathcal{S}$, then $\mathfrak{B} \in \mathcal{U}$.
2. A subclass \mathcal{U} of \mathcal{S} is said to be *preserved under extensions over \mathcal{S}* , abbreviated as \mathcal{U} is *PE over \mathcal{S}* , if for each structure $\mathfrak{A} \in \mathcal{U}$, if $\mathfrak{A} \subseteq \mathfrak{B}$ and $\mathfrak{B} \in \mathcal{S}$, then $\mathfrak{B} \in \mathcal{U}$.

If V and T are theories, then we say T is *PS modulo V* (resp. T is *PE modulo V*) if the class of models of $T \cup V$ is *PS* (resp. *PE*) over the class of models of V . For a sentence ϕ , we say ϕ is *PS modulo V* (resp. ϕ is *PE modulo V*) if the theory $\{\phi\}$ is *PS* (resp. *PE*) modulo V .

As an example, let $\tau = \{E\}$ be the vocabulary consisting of a single relation symbol E that is binary, and let \mathcal{S} be the class of all τ -structures in which E is interpreted as a symmetric binary relation. The class \mathcal{S} can be seen as the class of all undirected graphs. Let \mathcal{U}_1 be the subclass of \mathcal{S} consisting of all undirected graphs that are acyclic. Let \mathcal{U}_2 be the subclass of \mathcal{S} consisting of all undirected graphs that contain a triangle as a subgraph. It is easy to see that \mathcal{U}_1 is *PS* over \mathcal{S} , and \mathcal{U}_2 is *PE* over \mathcal{S} . Observe that \mathcal{S} is defined by the theory $V = \{\forall x \forall y (E(x, y) \rightarrow E(y, x))\}$. Let ψ_n be the universal sentence that asserts the absence of a cycle of length n as a subgraph. Then \mathcal{U}_1 is exactly the class of models in \mathcal{S} , of the theory $T = \{\psi_n \mid n \geq 3\}$, and \mathcal{U}_2 is exactly the class of models in \mathcal{S} , of the sentence $\phi = \neg\psi_3$. Whereby, T is *PS modulo V* , and ϕ is *PE modulo V* .

The following lemma establishes the duality between *PS* and *PE*.

Lemma 2.5.2 (*PS-PE duality*). *Let \mathcal{S} be a class of structures, \mathcal{U} be a subclass of \mathcal{S} and $\overline{\mathcal{U}}$ be the complement of \mathcal{U} in \mathcal{S} . Then \mathcal{U} is *PS* over \mathcal{S} iff $\overline{\mathcal{U}}$ is *PE* over \mathcal{S} . In particular, if \mathcal{S} is defined by a theory V , then a sentence ϕ is *PS modulo V* iff $\neg\phi$ is *PE modulo V* .*

The notion of a theory being *PS modulo V* or *PE modulo V* can be extended to theories with free variables in a natural manner. Given $n \in \mathbb{N}$, recall from Section 2.3 that τ_n is the vocabulary obtained by expanding τ with n fresh and distinct constants symbols c_1, \dots, c_n . Let $T(\bar{x})$ be an $\text{FO}(\tau)$ theory with free variables among $\bar{x} = (x_1, \dots, x_n)$, and let T' be the $\text{FO}(\tau_n)$ theory obtained by substituting c_i for the free occurrences of x_i in $T(\bar{x})$, for each $i \in \{1, \dots, n\}$. Given a theory V , we say $T(\bar{x})$ is *PS modulo V* if T' is *PS modulo V* , where V is treated as an $\text{FO}(\tau_n)$ theory. The notion $T(\bar{x})$ is *PE modulo V* is defined similarly.

2.5.1 The Łoś-Tarski preservation theorem

In the mid 1950s, Jerzy Łoś and Alfred Tarski provided syntactic characterizations of theories that are *PS* and theories that are *PE* via the following preservation theorem. This result (Theorem 3.2.2. in Chapter 3 of [12]) and its proof set the trend for various other preservation theorems to follow.

Theorem 2.5.3 (Łoś-Tarski, 1954-55). *Let $T(\bar{x})$ be a theory whose free variables are among \bar{x} . Given a theory V , each of the following is true.*

1. *$T(\bar{x})$ is PS modulo V iff $T(\bar{x})$ is equivalent modulo V to a theory $Y(\bar{x})$ of universal formulae, all of whose free variables are among \bar{x} . If $T(\bar{x})$ is a singleton, then so is $Y(\bar{x})$.*
2. *$T(\bar{x})$ is PE modulo V iff $T(\bar{x})$ is equivalent modulo V to a theory $Y(\bar{x})$ of existential formulae, all of whose free variables are among \bar{x} . If $T(\bar{x})$ is a singleton, then so is $Y(\bar{x})$.*

In the remainder of the thesis, if \mathcal{S} , as mentioned in the definitions above, is clear from context, then we skip mentioning its associated qualifier, namely, ‘over \mathcal{S} ’. Likewise, we skip mentioning ‘modulo V ’ when V is clear from context.

Chapter 3

New parameterized preservation properties

We fix a finite vocabulary τ in our discussion in this and in all the subsequent chapters of this part of the thesis. By formula, theory, and structure, we always mean respectively an $\text{FO}(\tau)$ formula, an $\text{FO}(\tau)$ theory and a τ -structure, unless explicitly stated otherwise.

3.1 Preservation under substructures modulo k -cruxes

Definition 3.1.1. Let \mathcal{S} be a class of structures and $k \in \mathbb{N}$. A subclass \mathcal{U} of \mathcal{S} is said to be *preserved under substructures modulo k -cruxes over \mathcal{S}* , abbreviated as \mathcal{U} is *PSC(k) over \mathcal{S}* , if for every structure $\mathfrak{A} \in \mathcal{U}$, there exists a subset C of the universe of \mathfrak{A} , of size at most k , such that if $\mathfrak{B} \subseteq \mathfrak{A}$, \mathfrak{B} contains C and $\mathfrak{B} \in \mathcal{S}$, then $\mathfrak{B} \in \mathcal{U}$. The set C is called a *k -crux of \mathfrak{A} w.r.t. \mathcal{U} over \mathcal{S}* . Any substructure \mathfrak{B} of \mathfrak{A} , that contains C is called a *substructure of \mathfrak{A} modulo the k -crux C* . Given theories V and T , we say *T is PSC(k) modulo V* , if the class of models of $T \cup V$ is PSC(k) over the class of models of V . For a sentence ϕ , we say *ϕ is PSC(k) modulo V* if the theory $\{\phi\}$ is PSC(k) modulo V .

Let $\mathcal{U}, \mathcal{S}, \mathfrak{A}, C, V, T$ and ϕ be as above. If \mathcal{S} is defined by V and \mathcal{U} is defined by T over \mathcal{S} , then we say *C is a k -crux of \mathfrak{A} w.r.t. T modulo V* . If \mathcal{U} is defined by ϕ over \mathcal{S} , then we say *C is a k -crux of \mathfrak{A} w.r.t. ϕ modulo V* . In many occasions in this thesis, the set C is the set of elements of a tuple \bar{a} of elements of \mathfrak{A} . Hence, we use the phrase *\bar{a} is a k -crux of \mathfrak{A} w.r.t. T modulo V* or *\bar{a} is a k -crux of \mathfrak{A} w.r.t. ϕ modulo V* to mean the corresponding statements with C in place of \bar{a} . As in Section 2.5, if any of $\mathcal{U}, \mathcal{S}, T, V$ or ϕ is clear from context, then we skip mentioning its

associated qualifier (viz., ‘w.r.t. \mathcal{U} ’, ‘over \mathcal{S} ’, ‘w.r.t. T ’, ‘modulo V ’ and ‘w.r.t. ϕ ’ respectively) in the definitions above.

Remark 3.1.2. Definition 3.1.1 is an adapted version of related definitions in [73] and [72]. The notion of ‘core’ in Definition 1 of [73] is exactly the notion of ‘crux’ defined above, where the underlying class \mathcal{U} in the definition above, is the class of all structures. We avoid using the word ‘core’ for a crux to prevent confusion with existing notions of cores in the literature [6, 70].

Given an $\text{FO}(\tau)$ theory $T(\bar{x})$ and an $\text{FO}(\tau)$ formula $\phi(\bar{x})$ each of whose free variables are among $\bar{x} = (x_1, \dots, x_n)$, we can define the notion of $T(\bar{x})$, resp. $\phi(\bar{x})$, being $PSC(k)$ modulo a theory V analogously to the notion of $T(\bar{x})$, resp. $\phi(\bar{x})$, being PS modulo V as defined in Section 2.5. Specifically, let c_1, \dots, c_n be the distinct constant symbols of $\tau_n \setminus \tau$, and let T' be the $\text{FO}(\tau_n)$ theory obtained by substituting c_i for the free occurrences of x_i in $T(\bar{x})$, for each $i \in \{1, \dots, n\}$. Then we say $T(\bar{x})$ is $PSC(k)$ modulo V if T' is $PSC(k)$ modulo V , where V is treated as an $\text{FO}(\tau_n)$ theory. The notion $\phi(\bar{x})$ is $PSC(k)$ modulo V is defined similarly.

Example 3.1.3. Let \mathcal{S} be the class of all undirected graphs. Given $k \in \mathbb{N}$, consider the class \mathcal{U}_k of all graphs of \mathcal{S} containing a cycle of length k as a subgraph. Clearly, for any graph G in \mathcal{U}_k , the vertices of any cycle of length k in G form a k -crux of G w.r.t. \mathcal{U}_k . Hence \mathcal{U}_k is $PSC(k)$. It is easy to see that \mathcal{U}_k is definable by an FO sentence, call it ϕ , whereby ϕ is $PSC(k)$.

Fix a class \mathcal{S} of structures. For properties P_1 and P_2 of subclasses of \mathcal{S} , we denote by $P_1 \Rightarrow P_2$ that any subclass of \mathcal{S} satisfying P_1 also satisfies P_2 . We denote by $P_1 \Leftrightarrow P_2$ that $P_1 \Rightarrow P_2$ and $P_2 \Rightarrow P_1$. It is now easy to check the following facts concerning the $PSC(k)$ subclasses of \mathcal{S} :

- (i) $PSC(0)$ coincides with the property of preservation under substructures, so $PSC(0) \Leftrightarrow PS$
- (ii) $PSC(l) \Rightarrow PSC(k)$ for $l \leq k$. If \mathcal{S} is any substructure-closed class of structures over a purely relational vocabulary (a vocabulary that contains only relation symbols), that contains infinitely many finite structures, then for each l , there exists $k > l$ and a $PSC(k)$ subclass \mathcal{U} of \mathcal{S} such that \mathcal{U} is not $PSC(l)$ over \mathcal{S} . This is seen as follows. Given l , let $k > l$ be such that there is some structure of size k in \mathcal{S} , and let ϕ_k be the sentence asserting that there are at least k elements in any model. Clearly ϕ_k is $PSC(k)$ over \mathcal{S} but not $PSC(l)$ over \mathcal{S} .

Define the property PSC of subclasses of \mathcal{S} as follows: A subclass \mathcal{U} of \mathcal{S} is PSC over \mathcal{S} if it is $PSC(k)$ over \mathcal{S} for some $k \in \mathbb{N}$. Notationally, $PSC \Leftrightarrow \bigvee_{k \geq 0} PSC(k)$. If \mathcal{S} is defined by a theory V , then the notions of ‘a sentence is PSC modulo V ’ and ‘a theory is PSC modulo V ’ are defined similarly as in Definition 3.1.1, and these notions are extended to formulae and

theories with free variables similarly as done above for $PSC(k)$. The implications mentioned in the previous paragraph show that PSC generalizes PS . If \mathcal{S} is any substructure-closed class of purely relational structures, that contains infinitely many finite structures, then the strict implications mentioned above show a strictly infinite hierarchy within PSC ; whence PSC provides a strict generalization of PS .

Suppose that \mathcal{S} is defined by a theory V . Given a Σ_2^0 sentence $\phi = \exists x_1 \dots \exists x_k \forall \bar{y} \varphi(x_1, \dots, x_k, \bar{y})$ and a structure \mathfrak{A} of \mathcal{S} such that $\mathfrak{A} \models \phi$, any set of *witnesses* in \mathfrak{A} of the existential quantifiers of ϕ , forms a k -crux of \mathfrak{A} . In particular, if a_1, \dots, a_k are witnesses in \mathfrak{A} , of the quantifiers associated with x_1, \dots, x_k (whence $\mathfrak{A} \models \forall \bar{y} \varphi(a_1, a_2, \dots, a_k, \bar{y})$), then given any substructure \mathfrak{B} of \mathfrak{A} containing a_1, \dots, a_k , the latter elements can again be chosen as witnesses in \mathfrak{B} , to make ϕ true in \mathfrak{B} . Therefore, ϕ is $PSC(k)$ (modulo V). It follows that Σ_2^0 formulae with k existential quantifiers are also $PSC(k)$ (modulo V).

Remark 3.1.4. Contrary to intuition, witnesses and k -cruxes cannot always be equated! Consider the sentence $\phi = \exists x \forall y E(x, y)$ and the structure $\mathfrak{A} = (\mathbb{N}, \leq)$, i.e. the natural numbers with the usual ordering. Let \mathcal{S} be the class of all structures. Clearly, ϕ is $PSC(1)$, $\mathfrak{A} \models \phi$ and the only witness of the existential quantifier of ϕ in \mathfrak{A} is the minimum element $0 \in \mathbb{N}$. In contrast, every singleton subset of \mathbb{N} is a 1-crux of \mathfrak{A} , since each substructure of \mathfrak{A} contains a minimum element under the induced order; this in turn is due to \mathbb{N} being well-ordered by \leq . This example shows that there can be models having many more (even infinitely more) cruxes than witnesses.

Since Σ_1^0 and Π_1^0 formulae are also Σ_2^0 formulae and the latter are PSC , the former are also PSC . However, Π_2^0 formulae are not necessarily PSC . Consider $\phi = \forall x \exists y E(x, y)$ and consider the model \mathfrak{A} of ϕ given by $\mathfrak{A} = (\mathbb{N}, E^{\mathfrak{A}} = \{(i, i+1) \mid i \in \mathbb{N}\})$. It is easy to check that no finite substructure of \mathfrak{A} models ϕ ; then \mathfrak{A} does not have any k -crux for any $k \in \mathbb{N}$, whence ϕ is not $PSC(k)$ for any k , and hence is not PSC .

3.2 Preservation under k -ary covered extensions

The classical notion of “extension of a structure” has a natural generalization to the notion of *extension of a collection of structures* as follows. A structure \mathfrak{A} is said to be an extension of a collection R of structures if for each $\mathfrak{B} \in R$, we have $\mathfrak{B} \subseteq \mathfrak{A}$. We now define a special kind of extensions of a collection of structures.

Definition 3.2.1. For $k \in \mathbb{N}$, a structure \mathfrak{A} is said to be a *k -ary covered extension* of a non-empty collection R of structures if (i) \mathfrak{A} is an extension of R , and (ii) for every subset C of the universe of \mathfrak{A} , of size at most k , there is a structure in R that contains C . We call R a *k -ary cover* of \mathfrak{A} .

Example 3.2.2. Let \mathfrak{A} be a graph on n vertices and let R be the collection of all r sized induced subgraphs of \mathfrak{A} , where $1 \leq r \leq n$. Then \mathfrak{A} is a k -ary covered extension of R for every k in $\{0, \dots, r\}$.

Remark 3.2.3. Note that a 0-ary covered extension of R is simply an extension of R . For $k > 0$, the universe of a k -ary covered extension of R is necessarily the union of the universes of the structures in R . However, different k -ary covered extensions of R can differ in the interpretation of predicates (if any) of arity greater than k . Note also that a k -ary covered extension of R is an l -ary covered extension of R for every $l \in \{0, \dots, k\}$.

Definition 3.2.4. Let \mathcal{S} be a class of structures and $k \in \mathbb{N}$. A subclass \mathcal{U} of \mathcal{S} is said to be *preserved under k -ary covered extensions* over \mathcal{S} , abbreviated as \mathcal{U} is *PCE(k)* over \mathcal{S} , if for every collection R of structures of \mathcal{U} , if \mathfrak{A} is a k -ary covered extension of R and $\mathfrak{A} \in \mathcal{S}$, then $\mathfrak{A} \in \mathcal{U}$. Given theories V and T , we say T is *PCE(k) modulo V* if the class of models of $T \cup V$ is *PCE(k)* over the class of models of V . For a sentence ϕ , we say ϕ is *PCE(k) modulo V* if the theory $\{\phi\}$ is *PCE(k) modulo V* .

As in the previous subsection, if any of \mathcal{S} or V is clear from context, then we skip mentioning its associated qualifier. Again as in the previous section, given a theory $T(\bar{x})$ and a formula $\phi(\bar{x})$ each of whose free variables is among \bar{x} , we can define the notions of ‘ $T(\bar{x})$ is *PCE(k) modulo V* ’ and ‘ $\phi(\bar{x})$ is *PCE(k) modulo V* ’ analogously to the corresponding notions in the context of *PSC(k)*.

The following lemma establishes the duality between *PSC(k)* and *PCE(k)*, generalizing the duality between *PS* and *PE* given by Lemma 2.5.2.

Lemma 3.2.5 (*PSC(k)-PCE(k) duality*). *Let \mathcal{S} be a class of structures, \mathcal{U} be a subclass of \mathcal{S} and $\overline{\mathcal{U}}$ be the complement of \mathcal{U} in \mathcal{S} . Then \mathcal{U} is *PSC(k)* over \mathcal{S} iff $\overline{\mathcal{U}}$ is *PCE(k)* over \mathcal{S} , for each $k \in \mathbb{N}$. In particular, if \mathcal{S} is defined by a theory V , then a sentence ϕ is *PSC(k) modulo V* iff $\neg\phi$ is *PCE(k) modulo V* .*

Proof. If: Suppose $\overline{\mathcal{U}}$ is $PCE(k)$ over \mathcal{S} but \mathcal{U} is not $PSC(k)$ over \mathcal{S} . Then there exists $\mathfrak{A} \in \mathcal{U}$ such that for every set C of at most k elements of \mathfrak{A} , there is a substructure \mathfrak{B}_C of \mathfrak{A} that (i) contains C , and (ii) belongs to $\mathcal{S} \setminus \mathcal{U}$, i.e. belongs to $\overline{\mathcal{U}}$. Then $R = \{\mathfrak{B}_C \mid C \text{ is a subset of } \mathfrak{A}, \text{ of size at most } k\}$ is a k -ary cover of \mathfrak{A} . Since $\overline{\mathcal{U}}$ is $PCE(k)$ over \mathcal{S} , it follows that $\mathfrak{A} \in \overline{\mathcal{U}}$ – a contradiction.

Only If: Suppose \mathcal{U} is $PSC(k)$ over \mathcal{S} but $\overline{\mathcal{U}}$ is not $PCE(k)$ over \mathcal{S} . Then there exists $\mathfrak{A} \in \mathcal{U}$ and a k -ary cover R of \mathfrak{A} such that every structure \mathfrak{B} of R belongs to $\overline{\mathcal{U}}$. Since \mathcal{U} is $PSC(k)$ over \mathcal{S} , there exists a k -crux C of \mathfrak{A} w.r.t. \mathcal{U} over \mathcal{S} . Consider the structure $\mathfrak{B}_C \in R$ that contains C – this exists since R is a k -ary cover of \mathfrak{A} . Then $\mathfrak{B}_C \in \mathcal{U}$ since C is a k -crux of \mathfrak{A} – a contradiction. \square

Fix a class \mathcal{S} of structures. Analogous to the notion of PSC , we define the notion of PCE as $PCE \Leftrightarrow \bigvee_{k \geq 0} PCE(k)$. The notions of a class, a sentence, a formula, a theory (without free variables) and a theory with free variables being PCE are defined analogously to corresponding notions for PSC . Then from the discussion in Section 3.1, and from Remark 3.2.3 and Lemma 3.2.5 above, we see that (i) $PCE(0) \Leftrightarrow PE$, (ii) $PCE(l) \Rightarrow PCE(k)$ for $l \leq k$, and (iii) a subclass \mathcal{U} of \mathcal{S} is PSC over \mathcal{S} iff the complement $\overline{\mathcal{U}}$ of \mathcal{U} in \mathcal{S} , is PCE over \mathcal{S} . Further, if \mathcal{S} is defined by a theory V , then all Π_2^0 formulae having at most k universal quantifiers are $PCE(k)$ (modulo V) and hence PCE , whereby all Σ_1^0 and Π_1^0 formulae are PCE as well. However Σ_2^0 formulae, in general, are not PCE since, as seen towards the end of Section 3.1, Π_2^0 formulae are, in general, not PSC .

In the next two chapters, we present characterizations of the $PSC(k)$ and $PCE(k)$ properties, and some natural variants of these. Our results and methods of proof are in general very different in the case of sentences vis-à-vis the case of theories. Hence, we deal with the two cases separately.

Chapter 4

Characterizations: the case of sentences

4.1 The generalized Łoś-Tarski theorem for sentences – $\text{GLT}(k)$

The central result of this section is as follows. This result, for the case of sentences, is called *the generalized Łoś-Tarski theorem for sentences at level k* , or simply *the generalized Łoś-Tarski theorem for sentences*, and is denoted as $\text{GLT}(k)$. Observe that for $k = 0$ below, we get exactly the Łoś-Tarski theorem for sentences.

Theorem 4.1.1. *Given a theory V , the following are true for each $k \in \mathbb{N}$.*

1. *A formula $\phi(\bar{x})$ is $\text{PSC}(k)$ modulo V iff $\phi(\bar{x})$ is equivalent modulo V to a Σ_2^0 formula whose free variables are among \bar{x} , and that has k existential quantifiers.*
2. *A formula $\phi(\bar{x})$ is $\text{PCE}(k)$ modulo V iff $\phi(\bar{x})$ is equivalent modulo V to a Π_2^0 formula whose free variables are among \bar{x} , and that has k universal quantifiers.*

Recall that $\text{PSC} \Leftrightarrow \bigvee_{k \geq 0} \text{PSC}(k)$ and $\text{PCE} \Leftrightarrow \bigvee_{k \geq 0} \text{PCE}(k)$. We then have the following corollary.

Corollary 4.1.2. *Given a theory V , the following are true.*

1. *A formula $\phi(\bar{x})$ is PSC modulo V iff $\phi(\bar{x})$ is equivalent modulo V to a Σ_2^0 formula whose free variables are among \bar{x} .*
2. *A formula $\phi(\bar{x})$ is PCE modulo V iff $\phi(\bar{x})$ is equivalent modulo V to a Π_2^0 formula whose free variables are among \bar{x} .*

The rest of this section is devoted to proving Theorem 4.1.1. We present two proofs of this result, one that uses λ -saturated structures (Section 4.1.1), and the other that uses ascending chains of structures (Section 4.1.2).

4.1.1 Proof of $\text{GLT}(k)$ using λ -saturated structures

Given theories T and V , we say that Γ is the *set of $\forall^k\exists^*$ consequences of T modulo V* if $\Gamma = \{\varphi \mid \varphi \text{ is a } \forall^k\exists^* \text{ sentence, and } T \text{ entails } \varphi \text{ modulo } V\}$. The following lemma is key to the proof.

Lemma 4.1.3. *Let V and T be consistent theories, and $k \in \mathbb{N}$. Let Γ be the set of $\forall^k\exists^*$ consequences of T modulo V . Then for all infinite cardinals μ , for every μ -saturated structure \mathfrak{A} that models V , we have that $\mathfrak{A} \models \Gamma$ iff there exists a k -ary cover R of \mathfrak{A} such that $\mathfrak{B} \models (V \cup T)$ for every $\mathfrak{B} \in R$.*

Proof. The ‘If’ direction is easy: for each $\mathfrak{B} \in R$, since $\mathfrak{B} \models (V \cup T)$, we have $\mathfrak{B} \models \varphi$ for each $\varphi \in \Gamma$. From the discussion towards the end of Section 3.2, any $\forall^k\exists^*$ sentence is $PCE(k)$ modulo V . Then since R is a k -ary cover of \mathfrak{A} , we have $\mathfrak{A} \models \varphi$ for each $\varphi \in \Gamma$.

For the ‘Only If’ direction, let the vocabulary of V and T be τ . We show that for every k -tuple \bar{a} of \mathfrak{A} , there is a substructure $\mathfrak{A}_{\bar{a}}$ of \mathfrak{A} containing (the elements of) \bar{a} such that $\mathfrak{A}_{\bar{a}} \models (V \cup T)$. Then the set $R = \{\mathfrak{A}_{\bar{a}} \mid \bar{a} \text{ is a } k\text{-tuple of } \mathfrak{A}\}$ forms the desired k -ary cover of \mathfrak{A} . To show the existence of $\mathfrak{A}_{\bar{a}}$, it suffices to show that there exists a τ -structure \mathfrak{B} such that (i) $|\mathfrak{B}| \leq \mu$, (ii) $\mathfrak{B} \models (V \cup T)$, and (iii) the Π_1^0 -type of \bar{a} in \mathfrak{A} , i.e. $\text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, \dots, x_k)$, is realized in \mathfrak{B} by some k -tuple, say \bar{b} . Then every Σ_1^0 sentence of $\text{FO}(\tau_k)$ true in (\mathfrak{B}, \bar{b}) is also true in (\mathfrak{A}, \bar{a}) . Since \mathfrak{A} is μ -saturated, we have by Proposition 2.4.2(1), that (\mathfrak{A}, \bar{a}) is also μ -saturated. There exists then, an isomorphic embedding $f : (\mathfrak{B}, \bar{b}) \rightarrow (\mathfrak{A}, \bar{a})$ by Proposition 2.4.2(5). Whereby the τ -reduct of the image of (\mathfrak{B}, \bar{b}) under f can serve as $\mathfrak{A}_{\bar{a}}$. The proof is therefore completed by showing the existence of \mathfrak{B} with the above properties.

Suppose $Z(x_1, \dots, x_k) = V \cup T \cup \text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, \dots, x_k)$ is inconsistent. By the compactness theorem, there is a finite subset of $Z(x_1, \dots, x_k)$ that is inconsistent. Since $\text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, \dots, x_k)$ is closed under taking finite conjunctions and since each of $\text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, \dots, x_k)$, V and T is consistent, there is a formula $\psi(x_1, \dots, x_k)$ in $\text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, \dots, x_k)$ such that $V \cup T \cup \{\psi(x_1, \dots, x_k)\}$ is inconsistent. In other words, $(V \cup T) \vdash \neg\psi(x_1, \dots, x_k)$. By \forall -introduction, we have $(V \cup T) \vdash \varphi$, where $\varphi = \forall x_1 \dots \forall x_k \neg\psi(x_1, \dots, x_k)$. Observe that φ is a $\forall^k\exists^*$ sentence; then by the definition of Γ , we have $\varphi \in \Gamma$, and hence $\mathfrak{A} \models \varphi$. Instantiating the k -tuple (x_1, \dots, x_k) as \bar{a} , we have $(\mathfrak{A}, \bar{a}) \models \neg\psi(x_1, \dots, x_k)$, contradicting the fact that $\psi(x_1, \dots, x_k) \in \text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, \dots, x_k)$. Then $Z(x_1, \dots, x_k)$ must be consistent. By the downward Löwenheim-Skolem theorem, there is a model (\mathfrak{B}, \bar{b}) of $Z(x_1, \dots, x_k)$ of power at most μ ; then \mathfrak{B} is as desired. \square

Proof of Theorem 4.1.1. We prove part (2) of Theorem 4.1.1. Part (1) of Theorem 4.1.1 follows from the duality of $PSC(k)$ and $PCE(k)$ given by Lemma 3.2.5. Also, we prove part (2) of Theorem 4.1.1 for the case of sentences; the result for formulae follows from definitions.

Suppose ϕ is equivalent modulo V to a $\forall^k\exists^*$ sentence φ . That φ is $PCE(k)$ modulo V follows from the discussion towards the end of Chapter 3. Whereby ϕ is $PCE(k)$ modulo V .

In the converse direction, suppose ϕ is $PCE(k)$ modulo V . If $V \cup \{\phi\}$ is unsatisfiable, we are trivially done. Otherwise, let Γ be the set of $\forall^k\exists^*$ consequences of $\{\phi\}$ modulo V . Then $(V \cup \{\phi\}) \vdash \Gamma$. We show below that $(V \cup \Gamma) \vdash \phi$, thereby showing that ϕ is equivalent to Γ modulo V . Then by the compactness theorem, we have ϕ is equivalent to a finite subset of Γ modulo V . Since a finite conjunction of $\forall^k\exists^*$ sentences is equivalent to a single $\forall^k\exists^*$ sentence, it follows that ϕ is equivalent to a $\forall^k\exists^*$ sentence, completing the proof.

Suppose $\mathfrak{A} \models (V \cup \Gamma)$. Consider a μ -saturated elementary extension \mathfrak{A}^+ of \mathfrak{A} , for some $\mu \geq \omega$ (\mathfrak{A}^+ exists by Proposition 2.4.2(3)). Then $\mathfrak{A}^+ \models (V \cup \Gamma)$. By Lemma 4.1.3, there exists a k -ary cover R of \mathfrak{A}^+ such that $\mathfrak{B} \models (V \cup \{\phi\})$ for every $\mathfrak{B} \in R$. Since ϕ is $PCE(k)$ modulo V , it follows that $\mathfrak{A}^+ \models \phi$. Then since $\mathfrak{A} \preceq \mathfrak{A}^+$, we have $\mathfrak{A} \models \phi$. \square

4.1.2 Proof of $\text{GLT}(k)$ using ascending chains of structures

We first define the notion of a k -ary cover of a structure \mathfrak{A} in an elementary extension of \mathfrak{A} . This notion generalizes the notion of k -ary cover seen earlier in Definition 3.2.1 – the latter corresponds to the notion in Definition 4.1.4 below, with \mathfrak{A}^+ being the same as \mathfrak{A} .

Definition 4.1.4. Let \mathfrak{A} be a structure and \mathfrak{A}^+ be an elementary extension of \mathfrak{A} . A non-empty collection R of substructures of \mathfrak{A}^+ is said to be a k -ary cover of \mathfrak{A} in \mathfrak{A}^+ if for every k -tuple \bar{a} of elements of \mathfrak{A} , there exists a structure in R containing \bar{a} .

The following lemma is key to the proof.

Lemma 4.1.5. Let V and T be consistent theories and $k \in \mathbb{N}$. Let Γ be the set of $\forall^k\exists^*$ consequences of T modulo V . Then for every structure \mathfrak{A} that models V , we have that $\mathfrak{A} \models \Gamma$ iff there exists an elementary extension \mathfrak{A}^+ of \mathfrak{A} and a k -ary cover R of \mathfrak{A} in \mathfrak{A}^+ such that $\mathfrak{B} \models (V \cup T)$ for every $\mathfrak{B} \in R$.

Proof. If: We show that $\mathfrak{A} \models \varphi$ for each sentence φ of Γ . Let $\varphi = \forall^k \bar{x} \psi(\bar{x})$ for a Σ_1^0 formula $\psi(\bar{x})$. Let \bar{a} be a k -tuple of \mathfrak{A} . Since R is a k -ary cover of \mathfrak{A} in \mathfrak{A}^+ , there exists $\mathfrak{B}_{\bar{a}} \in R$

such that $\mathfrak{B}_{\bar{a}}$ contains \bar{a} . Since $\mathfrak{B}_{\bar{a}} \models (V \cup T)$, we have $\mathfrak{B}_{\bar{a}} \models \Gamma$. Then $\mathfrak{B}_{\bar{a}} \models \varphi$ and hence $(\mathfrak{B}_{\bar{a}}, \bar{a}) \models \psi(\bar{x})$. Since $\psi(\bar{x})$ is a Σ_1^0 formula and $\mathfrak{B}_{\bar{a}} \subseteq \mathfrak{A}^+$, we have $(\mathfrak{A}^+, \bar{a}) \models \psi(\bar{x})$, whence $(\mathfrak{A}, \bar{a}) \models \psi(\bar{x})$ since $\mathfrak{A} \preceq \mathfrak{A}^+$. Since \bar{a} is arbitrary, $\mathfrak{A} \models \varphi$.

Only If: We have two cases here depending on whether \mathfrak{A} is finite or infinite. Before considering these cases, we present the following observation, call it \dagger . Let the vocabulary of \mathfrak{A} be τ .

(\dagger) Given an elementary extension \mathfrak{A}' of \mathfrak{A} and a k -tuple \bar{a} of \mathfrak{A} , there exist an elementary extension \mathfrak{A}'' of \mathfrak{A}' and a substructure \mathfrak{B} of \mathfrak{A}'' such that (i) \mathfrak{B} contains \bar{a} and (ii) $\mathfrak{B} \models (V \cup T)$.

This is seen as follows. Let $Z(\bar{x})$ be the theory given by $Z(\bar{x}) = V \cup T \cup \text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(\bar{x})$. We can show that $Z(\bar{x})$ is satisfiable by following the same argument as in the last paragraph of the proof of Lemma 4.1.3. Whereby if $(\mathfrak{D}, \bar{d}) \models Z(\bar{x})$, then every existential sentence that is true in (\mathfrak{D}, \bar{d}) is also true in (\mathfrak{A}, \bar{a}) , and hence in (\mathfrak{A}', \bar{a}) . Then by Lemma 2.3.4, there is an isomorphic embedding f of (\mathfrak{D}, \bar{d}) in an elementary extension $(\mathfrak{A}'', \bar{a})$ of (\mathfrak{A}, \bar{a}) . Taking \mathfrak{B} to be the τ -reduct of the image of (\mathfrak{D}, \bar{d}) under f , we see that \mathfrak{B} and \mathfrak{A}'' are indeed as desired.

We now consider the two cases mentioned above.

(1) \mathfrak{A} is finite: Given a k -tuple \bar{a} of \mathfrak{A} , by (\dagger), there exists an elementary extension \mathfrak{A}'' of \mathfrak{A} and a substructure $\mathfrak{B}_{\bar{a}}$ of \mathfrak{A}'' such that (i) $\mathfrak{B}_{\bar{a}}$ contains \bar{a} and (ii) $\mathfrak{B}_{\bar{a}} \models (V \cup T)$. Since \mathfrak{A} is finite, it follows from Lemma 2.3.1, that $\mathfrak{A}'' = \mathfrak{A}$. Whereby, taking $\mathfrak{A}^+ = \mathfrak{A}$ and $R = \{\mathfrak{B}_{\bar{a}} \mid \bar{a} \in U_{\mathfrak{A}}^k\}$, we see that \mathfrak{A}^+ and R are respectively indeed the desired elementary extension of \mathfrak{A} and k -ary cover of \mathfrak{A} in \mathfrak{A}^+ .

(2) \mathfrak{A} is infinite: The proof for this case is along the lines of the proof of the characterization of Π_2^0 sentences in terms of the property of preservation under unions of chains (see proof of Theorem 3.2.3 in Chapter 3 of [12]). Let λ be the successor cardinal of $|\mathfrak{A}|$ and $(\bar{a}_\kappa)_{\kappa < \lambda}$ be an enumeration of the k -tuples of \mathfrak{A} . For $\eta \leq \lambda$, given sequences $(\mathfrak{E}_\kappa)_{\kappa < \eta}$ and $(\mathfrak{F}_\kappa)_{\kappa < \eta}$ of structures, we say that $\mathcal{P}((\mathfrak{E}_\kappa)_{\kappa < \eta}, (\mathfrak{F}_\kappa)_{\kappa < \eta})$ is true iff $(\mathfrak{E}_\kappa)_{\kappa < \eta}$ is an ascending elementary chain and $\mathfrak{A} \preceq \mathfrak{E}_0$, and for each $\kappa < \eta$, we have (i) $\mathfrak{F}_\kappa \subseteq \mathfrak{E}_\kappa$ (ii) \mathfrak{F}_κ contains \bar{a}_κ and (iii) $\mathfrak{F}_\kappa \models (V \cup T)$. We then show the existence of sequences $(\mathfrak{A}_\kappa)_{\kappa < \lambda}$ and $(\mathfrak{B}_\kappa)_{\kappa < \lambda}$ of structures such that $\mathcal{P}((\mathfrak{A}_\kappa)_{\kappa < \lambda}, (\mathfrak{B}_\kappa)_{\kappa < \lambda})$ is true. Then by Theorem 2.3.5, taking $\mathfrak{A}^+ = \bigcup_{\kappa < \lambda} \mathfrak{A}_\kappa$ and $R = \{\mathfrak{B}_\kappa \mid \kappa < \lambda\}$, we see that \mathfrak{A}^+ and R are respectively indeed the desired elementary extension of \mathfrak{A} and k -ary cover of \mathfrak{A} in \mathfrak{A}^+ .

We construct the sequences $(\mathfrak{A}_\kappa)_{\kappa < \lambda}$ and $(\mathfrak{B}_\kappa)_{\kappa < \lambda}$ by constructing for each $\eta \leq \lambda$, the partial (initial) sequences $(\mathfrak{A}_\kappa)_{\kappa < \eta}$ and $(\mathfrak{B}_\kappa)_{\kappa < \eta}$ and showing that $\mathcal{P}((\mathfrak{A}_\kappa)_{\kappa < \eta}, (\mathfrak{B}_\kappa)_{\kappa < \eta})$ is true. We do this by (transfinite) induction on η . For the base case of $\eta = 1$, we see by (\dagger) above that if

$\mathfrak{A}' = \mathfrak{A}$, then there exists an elementary extension \mathfrak{A}'' of \mathfrak{A} and a substructure \mathfrak{B} of \mathfrak{A}'' such that (i) \mathfrak{B} contains \bar{a}_0 and (ii) $\mathfrak{B} \models (V \cup T)$. Then taking $\mathfrak{A}_0 = \mathfrak{A}''$ and $\mathfrak{B}_0 = \mathfrak{B}$, we see that $\mathcal{P}((\mathfrak{A}_0), (\mathfrak{B}_0))$ is true. As the induction hypothesis, assume that we have constructed sequences $(\mathfrak{A}_\kappa)_{\kappa < \eta}$ and $(\mathfrak{B}_\kappa)_{\kappa < \eta}$ such that $\mathcal{P}((\mathfrak{A}_\kappa)_{\kappa < \eta}, (\mathfrak{B}_\kappa)_{\kappa < \eta})$ is true. Then by Theorem 2.3.5, the structure $\mathfrak{A}' = \bigcup_{\kappa < \eta} \mathfrak{A}_\kappa$ is such that $\mathfrak{A} \preceq \mathfrak{A}'$. Then for the tuple \bar{a}_η of \mathfrak{A} , by (\dagger) , there exists an elementary extension \mathfrak{C} of \mathfrak{A}' and a substructure \mathfrak{D} of \mathfrak{C} such that (i) \mathfrak{D} contains \bar{a}_η and (ii) $\mathfrak{D} \models (V \cup T)$. Then taking $\mathfrak{A}_\eta = \mathfrak{C}$ and $\mathfrak{B}_\eta = \mathfrak{D}$, and letting μ be the successor ordinal of η , we see that $\mathcal{P}((\mathfrak{A}_\kappa)_{\kappa < \mu}, (\mathfrak{B}_\kappa)_{\kappa < \mu})$ is indeed true, completing the induction. \square

Proof of Theorem 4.1.1. We prove part (2) of Theorem 4.1.1. Part (1) of Theorem 4.1.1 follows from the duality of $PSC(k)$ and $PCE(k)$ given by Lemma 3.2.5. We prove part (2) of Theorem 4.1.1 for the case of sentences; the result for formulae follows. The ‘If’ direction of part (2) of Theorem 4.1.1 is proved exactly as the proof of this part of Theorem 4.1.1, as presented in the Section 4.1.1. We hence prove the ‘Only if’ direction below.

Suppose ϕ is $PCE(k)$ modulo V . If $V \cup \{\phi\}$ is unsatisfiable, we are trivially done. Otherwise, let Γ be the set of $\forall^k \exists^*$ consequences of $\{\phi\}$ modulo V . Then $(V \cup \{\phi\}) \vdash \Gamma$. We show below that $(V \cup \Gamma) \vdash \phi$, thereby showing that ϕ is equivalent to Γ modulo V . Suppose $\mathfrak{A} \models (V \cup \Gamma)$. Consider the sequence $(\mathfrak{A}_i)_{i \geq 0}$ of structures and the sequence $(R_i)_{i \geq 0}$ of collections of structures with the following properties.

1. $(\mathfrak{A}_i)_{i \geq 0}$ is an ascending elementary chain such that $\mathfrak{A} \preceq \mathfrak{A}_0$ (whereby $\mathfrak{A}_i \models (V \cup \Gamma)$ for each $i \geq 0$) and for each $i \geq 0$, \mathfrak{A}_{i+1} is the elementary extension of \mathfrak{A}_i as given by Lemma 4.1.5.
2. For each $i \geq 0$, R_i is the k -ary cover of \mathfrak{A}_i in \mathfrak{A}_{i+1} as given by Lemma 4.1.5.

Consider the structure $\mathfrak{A}^+ = \bigcup_{i \geq 0} \mathfrak{A}_i$. Consider any k -tuple \bar{a} of \mathfrak{A}^+ ; it is clear that there must exist $j \geq 0$ such \bar{a} is contained in \mathfrak{A}_j . Then there exists a structure $\mathfrak{B}_{\bar{a}} \in R_j$ such that (i) $\mathfrak{B}_{\bar{a}}$ contains \bar{a} and (ii) $\mathfrak{B}_{\bar{a}} \models (V \cup \{\phi\})$. Since $\mathfrak{B}_{\bar{a}} \in R_j$, we have $\mathfrak{B}_{\bar{a}} \subseteq \mathfrak{A}_{j+1}$ and since $\mathfrak{A}_{j+1} \preceq \mathfrak{A}^+$ (by Theorem 2.3.5), we have $\mathfrak{B}_{\bar{a}} \subseteq \mathfrak{A}^+$. Then $R = \{\mathfrak{B}_{\bar{a}} \mid \bar{a} \text{ is a } k\text{-tuple from } \mathfrak{A}^+\}$ is a k -ary cover of \mathfrak{A}^+ (or equivalently, a k -ary cover of \mathfrak{A}^+ in \mathfrak{A}^+) such that $\mathfrak{B} \models (V \cup \{\phi\})$ for each $\mathfrak{B} \in R$. Since ϕ is $PCE(k)$ modulo V , it follows that $\mathfrak{A}^+ \models \phi$. Then since $\mathfrak{A} \preceq \mathfrak{A}^+$, we have that $\mathfrak{A} \models \phi$, completing the proof. \square

4.2 Variants of our properties and their characterizations

In this section, we present natural generalizations of the $PSC(k)$ and $PCE(k)$ properties in which, rather than insisting on bounded sized cruxes and bounded arity covers, we allow cruxes of sizes, and covers of arities, less than λ , where λ is an infinite cardinal. We first define the notion of λ -ary covered extensions.

Definition 4.2.1. Given an infinite cardinal λ , a structure \mathfrak{A} is called a λ -ary covered extension of a collection R of structures if (i) \mathfrak{A} is an extension of R (ii) for each subset C of the universe of \mathfrak{A} , of size less than λ , there is a structure in R containing C . We call R a λ -ary cover of \mathfrak{A} .

Observe that in the definition above, \mathfrak{A} must be unique such since all relation symbols and function symbols have finite arity.

Definition 4.2.2. Let \mathcal{S} be a class of structures and \mathcal{U} be a subclass of \mathcal{S} .

1. We say \mathcal{U} is *preserved under substructures modulo λ -cruxes over \mathcal{S}* , abbreviated \mathcal{U} is $PSC(\lambda)$ over \mathcal{S} , if for each structure $\mathfrak{A} \in \mathcal{U}$, there is a subset C of the universe of \mathfrak{A} , of size less than λ , such that, if $\mathfrak{B} \subseteq \mathfrak{A}$, \mathfrak{B} contains C and $\mathfrak{B} \in \mathcal{S}$, then $\mathfrak{B} \in \mathcal{U}$. The set C is called an λ -crux of \mathfrak{A} w.r.t. \mathcal{U} over \mathcal{S} .
2. We say \mathcal{U} is *preserved under λ -ary covered extensions over \mathcal{S}* , abbreviated \mathcal{U} is $PCE(\lambda)$ over \mathcal{S} , if for every collection R of structures of \mathcal{U} , if \mathfrak{A} is an λ -ary covered extension of R and $\mathfrak{A} \in \mathcal{S}$, then $\mathfrak{A} \in \mathcal{U}$.

It is easy to see that given classes \mathcal{U} and \mathcal{S} , and infinite cardinals λ and μ such that $\lambda \leq \mu$, if \mathcal{U} is $PSC(\lambda)$ (resp. $PCE(\lambda)$) over \mathcal{S} , then \mathcal{U} is $PSC(\mu)$ (resp. $PCE(\mu)$) over \mathcal{S} .

If $\phi(\bar{x})$ and $T(\bar{x})$ are respectively a formula and a theory with free variables \bar{x} , then given a theory V , the notions of ‘ $\phi(\bar{x})$ is $PSC(\lambda)$ ($PCE(\lambda)$) modulo V ’ and ‘ $T(\bar{x})$ is $PSC(\lambda)$ ($PCE(\lambda)$) modulo V ’ are defined similarly as the corresponding notions for $PSC(k)$ and $PCE(k)$.

Analogous to Lemma 3.2.5, Lemma 4.1.3, Lemma 4.1.5 and Theorem 4.1.1, we have the following results for $PSC(\lambda)$ and $PCE(\lambda)$. The proofs are similar to the corresponding results for $PSC(k)$ and $PCE(k)$ and are hence skipped.

Lemma 4.2.3 ($PSC(\lambda)$ - $PCE(\lambda)$ duality). *Let \mathcal{S} be a class of structures, \mathcal{U} be a subclass of \mathcal{S} and $\bar{\mathcal{U}}$ be the complement of \mathcal{U} in \mathcal{S} . Then \mathcal{U} is $PSC(\lambda)$ over \mathcal{S} iff $\bar{\mathcal{U}}$ is $PCE(\lambda)$ over \mathcal{S} . In particular, if \mathcal{S} is defined by a theory V , then a sentence ϕ is $PSC(\lambda)$ modulo V iff $\neg\phi$ is $PCE(\lambda)$ modulo V .*

Lemma 4.2.4. *Let V and T be consistent theories, and let Γ be the set of Π_2^0 consequences of T modulo V . Then for all infinite cardinals λ and μ , and for every μ -saturated structure \mathfrak{A} that models V , we have that $\mathfrak{A} \models \Gamma$ iff there exists a λ -ary cover R of \mathfrak{A} such that $\mathfrak{B} \models (V \cup T)$ for every $\mathfrak{B} \in R$.*

Lemma 4.2.5. *Let V and T be consistent theories, and let Γ be the set of Π_2^0 consequences of T modulo V . Then for all infinite cardinals λ , and for every structure \mathfrak{A} that models V , we have that $\mathfrak{A} \models \Gamma$ iff there exists an elementary extension \mathfrak{A}^+ of \mathfrak{A} and a λ -ary cover R of \mathfrak{A} in \mathfrak{A}^+ such that $\mathfrak{B} \models (V \cup T)$ for every $\mathfrak{B} \in R$.*

Theorem 4.2.6. *Given a theory V , the following hold for each infinite cardinal λ .*

1. *A formula $\phi(\bar{x})$ is $PSC(\lambda)$ modulo V iff $\phi(\bar{x})$ is equivalent modulo V to a Σ_2^0 formula having free variables \bar{x} .*
2. *A formula $\phi(\bar{x})$ is $PCE(\lambda)$ modulo V iff $\phi(\bar{x})$ is equivalent modulo V to a Π_2^0 formula having free variables \bar{x} .*

The above theorem implies the following result that is not obvious from the definitions of the properties concerned.

Corollary 4.2.7. *For every infinite cardinal λ , a sentence is $PSC(\lambda)$ (resp. $PCE(\lambda)$) modulo a theory V iff it is PSC (resp. PCE) modulo V .*

The above characterizations, along with the characterizations in Section 4.1, are depicted pictorially below.

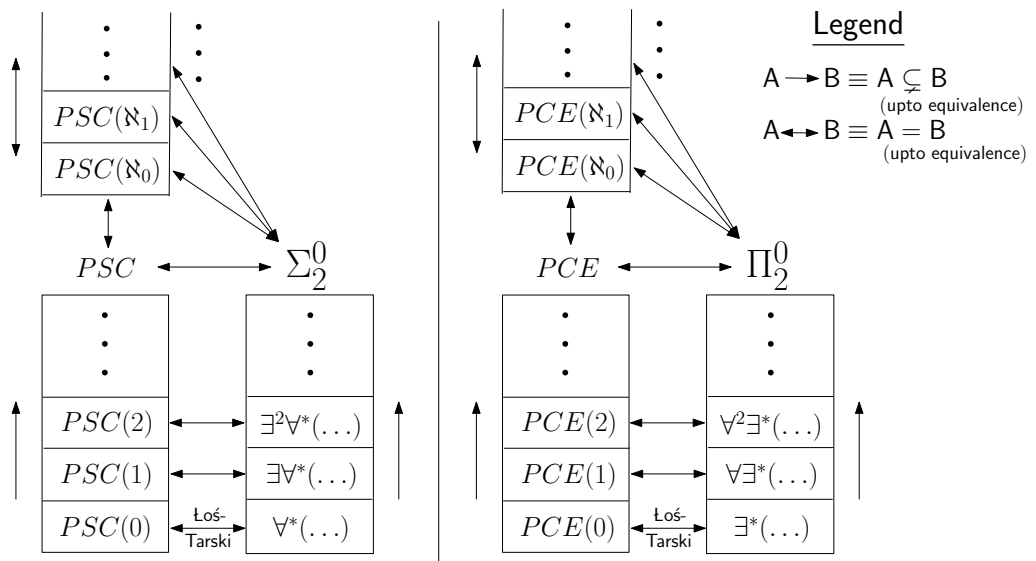


Figure 4.1: Characterizations of $PSC(k)$, $PCE(k)$, $PSC(\lambda)$ and $PCE(\lambda)$ sentences

4.2.1 Applications: new proofs of inexpressibility results in FO

Typical proofs of inexpressibility results in FO are either via compactness theorem, or Ehrenfeucht-Fr  iss   games or locality arguments. We present a new approach to proving inexpressibility results, using our results. We illustrate this approach via the example presented below. Below, the underlying class \mathcal{S} of graphs is the class of all undirected graphs.

Consider the subclass \mathcal{U} of \mathcal{S} consisting of graphs that contain a cycle as a subgraph. It is easy to see that in any graph $G \in \mathcal{U}$, the vertices of any cycle form an \aleph_0 -crux of G . Then \mathcal{U} is $PSC(\aleph_0)$. If \mathcal{U} were definable by an FO sentence, say φ , then φ is $PSC(\aleph_0)$. By Corollary 4.2.7, it follows that φ is $PSC(k)$ for some $k \in \mathbb{N}$. Now consider the cycle graph G of length $k + 1$; clearly G models φ . No proper induced subgraph of G is a cycle, whence G contains no k -crux at all. This contradicts the earlier inference that φ is $PSC(k)$. Thus \mathcal{U} is not definable by any FO sentence.

A short report containing more examples of inexpressibility results proven using our preservation theorems can be found at [74]. These examples include connectedness, bipartiteness, caterpillars, etc. Note that the notion of ‘core’ in [74] is exactly what we mean by a ‘crux’ in this thesis.

4.3 An uncomputability result

Corollary 4.2.7 tells that given a sentence ϕ that is $PSC(\lambda)$ (resp. $PCE(\lambda)$) modulo a theory V , there exists $k \in \mathbb{N}$ such that ϕ is $PSC(k)$ (resp. $PCE(k)$) modulo V . This raises the question: is k computable? The following proposition answers the aforesaid question in the negative for $PSC(\aleph_0)$ (resp. $PCE(\aleph_0)$), and hence for $PSC(\lambda)$ (resp. $PCE(\lambda)$) for each (infinite cardinal) λ . Below, a *relational* sentence is a sentence over a vocabulary that does not contain any function symbols. Let the length of a sentence ϕ be denoted by $|\phi|$.

Proposition 4.3.1. *Let V be the empty theory. For every recursive function $\nu : \mathbb{N} \rightarrow \mathbb{N}$, the following are true:*

1. *There is a relational Π_2^0 sentence ϕ that is $PSC(\aleph_0)$ modulo V but that is not $PSC(k)$ modulo V for any $k \leq \nu(|\phi|)$.*
2. *There is a relational Σ_2^0 sentence ϕ that is $PCE(\aleph_0)$ modulo V but that is not $PCE(k)$ modulo V for any $k \leq \nu(|\phi|)$.*

Towards the proof of the above proposition, we first present a recent unpublished result of Rossman [71].

Theorem 4.3.2 (Rossman, 2012). *Let V be the empty theory. For every recursive function $\nu : \mathbb{N} \rightarrow \mathbb{N}$, there exists a relational Σ_2^0 sentence ϕ that is PS modulo V , and for which every equivalent Π_1^0 sentence has length at least $\nu(|\phi|) + 1$.*

Theorem 4.3.2 gives a non-recursive lower bound on the length of Π_1^0 sentences equivalent to sentences that are PS (in terms of the lengths of the latter sentences). This strengthens the non-elementary lower bound proved in [18].

Corollary 4.3.3. *Let V be the empty theory. For every recursive function $\nu : \mathbb{N} \rightarrow \mathbb{N}$, there exists a relational Σ_2^0 sentence ϕ that is PS modulo V , and for which every equivalent Π_1^0 sentence has at least $\nu(|\phi|) + 1$ universal variables.*

Proof. We show below that there is a monotone recursive function $\rho : \mathbb{N} \rightarrow \mathbb{N}$ such that if ξ is a Π_1^0 sentence with n variables, then the shortest (in terms of length) Π_1^0 sentence equivalent to ξ has length at most $\rho(n)$. That would prove this corollary as follows. Suppose there is a recursive function $\nu : \mathbb{N} \rightarrow \mathbb{N}$ such that for each relational Σ_2^0 sentence ψ that is PS modulo V , there is an equivalent Π_1^0 sentence having at most $\nu(|\psi|)$ universal variables. Then consider the recursive function $\theta : \mathbb{N} \rightarrow \mathbb{N}$ given by $\theta(n) = \rho(\nu(n))$ and let ϕ be the relational Σ_2^0 sentence given by Theorem 4.3.2 for the function θ . Then ϕ is PS modulo V and the shortest Π_1^0 sentence equivalent to ϕ has length $> \theta(|\phi|)$. By the assumption about ν above, there is a Π_1^0 sentence equivalent to ϕ having at most $\nu(|\phi|)$ universal variables. Whence there is a Π_1^0 sentence equivalent to ϕ whose length is at most $\rho(\nu(|\phi|)) = \theta(|\phi|)$ – a contradiction.

Let ξ be a universal sentence given by $\xi = \forall^n \bar{z} \beta(\bar{z})$. Let the vocabulary of ξ be τ and the maximum arity of any predicate of τ be q . Then the number k of atomic formulae of τ having variables from \bar{z} is at most $|\tau| \cdot n^q$. It follows that the length r of the disjunctive normal form, say α , of β satisfies $r \leq (d \cdot k \cdot 2^k)$ for some constant $d \geq 1$. Then ξ is equivalent to the sentence $\gamma = \forall^n \bar{z} \alpha(\bar{z})$; the size of γ is at most $e \cdot (n + r)$ for some constant $e \geq 1$. Since k and r are bounded by monotone recursive functions of n , so is the length of γ . \square

Proof of Proposition 4.3.1. We give the proof for part (1). The negation of the sentence ϕ showing part (1) proves part (2). Also, we omit the mention of V for the sake of readability.

Suppose there is a recursive function $\nu : \mathbb{N} \rightarrow \mathbb{N}$ such that if ξ is a relational Π_2^0 sentence that is $PSC(\aleph_0)$, then ξ is $PSC(k)$ for some $k \leq \nu(|\xi|)$. In other words, for ξ as mentioned, every

model of ξ has a crux of size at most $\nu(|\xi|)$. Consider the recursive function $\rho : \mathbb{N} \rightarrow \mathbb{N}$ given by $\rho(n) = \nu(n + 1)$. Then, for the function ρ , consider the relational Σ_2^0 sentence ϕ given by Corollary 4.3.3. The sentence ϕ is *PS* and every Π_1^0 sentence equivalent to it has $> \rho(|\phi|)$ number of universal variables. Now the Π_2^0 sentence ψ given by $\psi = \neg\phi$ is equivalent to a Σ_1^0 sentence. Since Σ_1^0 sentences are *PSC*, and hence *PSC*(\aleph_0), it follows that ψ is *PSC*(\aleph_0). Now, by our assumption about ν above, every model of ψ has a crux of size at most $\nu(|\psi|) = \nu(|\phi| + 1) = \rho(|\phi|)$. Then all minimal models of ψ have size at most $\rho(|\phi|) + q$, where q is the number of constant symbols in the vocabulary of ϕ . Using the fact that ψ is preserved under extensions, it is easy to construct a Σ_1^0 sentence having $\rho(|\phi|)$ number of existential variables, that is equivalent to ψ . Whereby ϕ is equivalent to a Π_1^0 sentence having $\rho(|\phi|)$ number of universal variables – a contradiction. \square

Chapter 5

Characterizations: the case of theories

5.1 Characterizations of the extensional properties

The central result of this section is as below.

Theorem 5.1.1. *Given a theory V , the following hold for each $k \in \mathbb{N}$ and each $\lambda \geq \aleph_0$:*

1. *A theory $T(\bar{x})$ is $PCE(k)$ modulo V iff $T(\bar{x})$ is equivalent modulo V to a theory of Π_2^0 formulae, all of whose free variables are among \bar{x} and all of which have k universal quantifiers.*
2. *A theory $T(\bar{x})$ is $PCE(\lambda)$ modulo V iff $T(\bar{x})$ is equivalent modulo V to a theory of Π_2^0 formulae, all of whose free variables are among \bar{x} .*

The proofs of part (1) and part (2) of the above result are respectively, nearly identical to the proofs of Theorem 4.1.1(2) and Theorem 4.2.6(2) – we just consider theories instead of sentences in the latter proofs and use the following lemma that is straightforward.

Lemma 5.1.2. *Let \mathcal{S} be a class of structures, k a natural number and λ an infinite cardinal. For an index set I , let $\{\mathcal{U}_i \mid i \in I\}$ be a collection of subclasses of \mathcal{S} such that \mathcal{U}_i is $PCE(k)$, resp. $PCE(\lambda)$, over \mathcal{S} , for each $i \in I$. Then $\bigcap_{i \in I} \mathcal{U}_i$ is $PCE(k)$, resp. $PCE(\lambda)$, over \mathcal{S} .*

Remark 5.1.3. By considering singleton theories in Theorem 5.1.1, and using compactness theorem and the fact that a finite conjunction of $\forall^k \exists^*$ sentences, respectively Π_2^0 sentences, is also a $\forall^k \exists^*$ sentence, respectively a Π_2^0 sentence, we get Theorem 4.1.1(2) and Theorem 4.2.6(2).

The following proposition reveals an important difference between considering the properties of $PCE(k)$ and $PCE(\lambda)$ in the context of theories, vis-à-vis considering these properties in

the context of sentences. Specifically, in contrast to Corollary 4.2.7, it turns out that $PCE(\lambda)$ theories are more general than PCE theories.

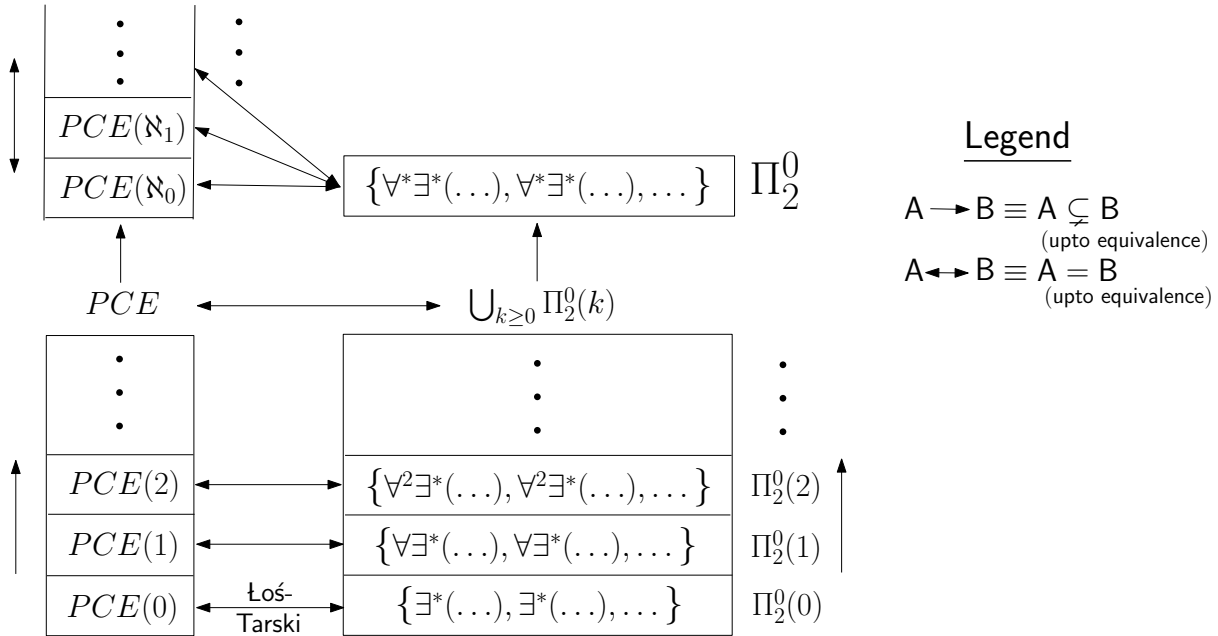
Proposition 5.1.4. *Let λ be an infinite cardinal.*

1. *A theory is $PCE(\lambda)$ modulo a theory V iff it is $PCE(\aleph_0)$ modulo V .*
2. *There are theories T and V such that T is $PCE(\aleph_0)$ modulo V , and hence $PCE(\lambda)$ modulo V , but T is not PCE modulo V .*

Proof. Part (1) follows easily from Theorem 5.1.1(2). We prove part (2) below.

Let V be the theory defining the class of all undirected graphs. Let T be a Π_1^0 theory over graphs asserting that there is no cycle of length k for any $k \in \mathbb{N}$. Then T defines the class \mathcal{U} of all acyclic graphs, and is $PCE(\aleph_0)$ modulo V by Theorem 5.1.1(2). Suppose T is PCE modulo V , whence T is $PCE(k)$ modulo V for some $k \in \mathbb{N}$. Then \mathcal{U} is $PCE(k)$ modulo the class of models of V . By Lemma 3.2.5, $\overline{\mathcal{U}}$ (the complement of \mathcal{U}) is $PSC(k)$ modulo the class of models of V . Now consider a cycle G of length $k + 1$. Clearly, G is in $\overline{\mathcal{U}}$ but every proper substructure of G is in \mathcal{U} . This contradicts our earlier inference that $\overline{\mathcal{U}}$ is $PSC(k)$ modulo the class of models of V . \square

The characterizations of this section are depicted pictorially below.



5.2 Characterizations of the substructural properties

The central results of this section are as follows.

Theorem 5.2.1. *Let V be a given theory, $k \in \mathbb{N}$ and $\lambda > \aleph_0$.*

1. *A theory $T(\bar{x})$ is $PSC(\lambda)$ modulo V iff $T(\bar{x})$ is equivalent modulo V to a theory of Σ_2^0 formulae, all of whose free variables are among \bar{x} .*
2. *If a theory $T(\bar{x})$ is $PSC(\aleph_0)$ modulo V , then $T(\bar{x})$ is equivalent modulo V to a theory of Σ_2^0 formulae, all of whose free variables are among \bar{x} . The same consequent (therefore) holds if $T(\bar{x})$ is $PSC(k)$ modulo V . The converses of these implications are not true. There exist theories T and V such that (i) each sentence of T is a Σ_2^0 sentence having exactly one existential quantifier, and (ii) T is not $PSC(\aleph_0)$ modulo V , and hence not $PSC(k)$ modulo V .*

Since Theorem 4.1.1(1) shows that a $PSC(k)$ sentence is always equivalent to an $\exists^k \forall^*$ sentence, it is natural to ask if a $PSC(k)$ theory is always equivalent to a theory of $\exists^k \forall^*$ sentences. We give an affirmative answer to this question, conditioned on a hypothesis that we present below, and thereby provide a (conditional) refinement of Theorem 5.2.1(2).

Hypothesis 5.2.2. *Given theories V and $T(\bar{x})$, and $k \in \mathbb{N}$ such that $T(\bar{x})$ is $PSC(k)$ modulo V , it is the case that for each model (\mathfrak{A}, \bar{a}) of $T(\bar{x})$, there exists a k -crux \bar{b} of (\mathfrak{A}, \bar{a}) (w.r.t. $T(\bar{x})$ modulo V) such that \bar{b} is also a k -crux (w.r.t. $T(\bar{x})$ modulo V) of a μ -saturated elementary extension of (\mathfrak{A}, \bar{a}) , for some $\mu \geq \omega$.*

Theorem 5.2.3. *Given theories V and $T(\bar{x})$, suppose $T(\bar{x})$ is $PSC(k)$ modulo V for a given $k \in \mathbb{N}$. Then assuming Hypothesis 5.2.2, $T(\bar{x})$ is equivalent modulo V to a theory of Σ_2^0 formulae, all of whose free variables are among \bar{x} , and all of which have k existential quantifiers.*

The approach of ‘dualizing’ adopted in proving Theorem 4.1.1(1) cannot work for characterizing theories that are $PSC(k)$ or $PSC(\lambda)$ since the negation of an FO theory might, in general, not be equivalent to any FO theory. We therefore present in this section, altogether different approaches to proving the above results. While we show that Σ_2^0 theories characterize $PSC(\lambda)$ theories for $\lambda > \aleph_0$, it is unclear at present what syntactic fragments of FO theories serve to characterize $PSC(k)$ and $PSC(\aleph_0)$ theories. However, Theorem 5.2.1(2) shows that these syntactic fragments must be semantically contained inside the class of Σ_2^0 theories, and Theorem 5.2.3 shows that under Hypothesis 5.2.2, any syntactic fragment that characterizes $PSC(k)$ theories must be semantically contained inside the class of theories of $\exists^k \forall^*$ sentences.

The remainder of this section is entirely devoted to proving the results above. In the next two sections, we present the proofs of Theorem 5.2.1 and Theorem 5.2.3, and also show that Hypothesis 5.2.2 is indeed well-motivated.

Proof of Theorem 5.2.1

We first present the proof of part (2) of Theorem 5.2.1, assuming part (1) of Theorem 5.2.1.

Proof of Theorem 5.2.1(2). Since a theory $T(\bar{x})$ that is $PSC(k)$ modulo V or $PSC(\aleph_0)$ modulo V is also $PSC(\lambda)$ modulo V for $\lambda > \aleph_0$, it follows from Theorem 5.2.1(1) that $T(\bar{x})$ is equivalent modulo V to a theory of Σ_2^0 formulae, all of whose free variables are among \bar{x} . We show below that the converse is not true.

Let $V = \{\forall x \forall y (E(x, y) \rightarrow E(y, x))\}$ be the theory that defines exactly all undirected graphs. For $n \geq 1$, let $\varphi_n(x)$ be a formula asserting that x is not a part of a cycle of length n . Explicitly, $\varphi_1(x) = \neg E(x, x)$ and for $n \geq 1$, we have $\varphi_{n+1}(x) = \neg \exists z_1 \dots \exists z_n ((\bigwedge_{1 \leq i < j \leq n} z_i \neq z_j) \wedge (\bigwedge_{i=1}^{i=n} (x \neq z_i)) \wedge E(x, z_1) \wedge E(z_n, x) \wedge \bigwedge_{i=1}^{i=n-1} E(z_i, z_{i+1}))$. Consider $\chi_n(x) = \bigwedge_{i=1}^{i=n} \varphi_i(x)$ which asserts that x is not a part of any cycle of length $\leq n$. Observe that $\chi_n(x)$ is equivalent to a universal formula. Also, if $m \leq n$, then $\chi_n(x) \rightarrow \chi_m(x)$.

Now consider the theory $T = \{\psi_n \mid n \geq 1\}$, where $\psi_n = \exists x \chi_n(x)$. Each sentence of T is a Σ_2^0 sentence having only one existential quantifier. We show that T is not $PSC(\aleph_0)$ modulo V .

Consider the infinite graph G given by $G = \bigsqcup_{i \geq 3} C_i$ where C_i is the cycle graph of length i and \bigsqcup denotes disjoint union. Any vertex x of C_i satisfies $\chi_j(x)$ in G , for $j < i$. Then $G \models T$. Now consider any finite set S of vertices of G . Let r be the highest index such that some vertex in S is in the cycle C_r . Consider the subgraph G_1 of G induced by the vertices of all the cycles in G of length $\leq r$. Then no vertex x of G_1 satisfies $\chi_l(x)$ for $l > r$. Then $G_1 \not\models T$, whence S cannot be a k -crux of G w.r.t. T modulo V , for any $k \geq |S|$. Since S is an arbitrary finite subset of G , we conclude that G has no k -crux w.r.t. T modulo V , for any $k \in \mathbb{N}$; in other words, G has no \aleph_0 -crux. Then T is not $PSC(\aleph_0)$ modulo V . \square

Towards the proof of part (1) of Theorem 5.2.1, we recall the notion of *sandwiches* as defined by Keisler in [45]. We say that a triple $(\mathfrak{A}, \mathfrak{B}, \mathfrak{C})$ of structures is a *sandwich* if $\mathfrak{A} \preceq \mathfrak{C}$ and $\mathfrak{A} \subseteq \mathfrak{B} \subseteq \mathfrak{C}$. Given structures \mathfrak{A} and \mathfrak{B} , we say that \mathfrak{B} is *sandwiched* by \mathfrak{A} if there exist structures \mathfrak{A}' and \mathfrak{B}' such that (i) $\mathfrak{B} \preceq \mathfrak{B}'$ and (ii) $(\mathfrak{A}, \mathfrak{B}', \mathfrak{A}')$ is a sandwich. Given theories V and T , we say T is *preserved under sandwiches by models of T modulo V* if for each model \mathfrak{A}

of $V \cup T$, if \mathfrak{B} is sandwiched by \mathfrak{A} and \mathfrak{B} models V , then \mathfrak{B} models T . The following theorem of Keisler (Corollary 5.2 of [45]) gives a syntactic characterization of the aforesaid preservation property in terms of Σ_2^0 theories.

Theorem 5.2.4 (Keisler, 1960). *Let V and T be theories. Then T is preserved under sandwiches by models of T modulo V iff T is equivalent modulo V to a theory of Σ_2^0 sentences.*

To prove the ‘Only if’ direction of Theorem 5.2.1(1) therefore, it just suffices to show that if T is a theory that is $PSC(\lambda)$ modulo V , then T is preserved under sandwiches by models of T modulo V . To do this, we first prove the following lemmas.

Lemma 5.2.5 (Sandwich by saturated structures). *Let \mathfrak{A}_1 and \mathfrak{B}_1 be structures such that \mathfrak{B}_1 is sandwiched by \mathfrak{A}_1 . Then for each $\mu \geq \omega$, for every μ -saturated elementary extension \mathfrak{A} of \mathfrak{A}_1 , there exists a structure \mathfrak{B} isomorphic to \mathfrak{B}_1 such that \mathfrak{B} is sandwiched by \mathfrak{A} .*

Lemma 5.2.6 (Preservation under sandwich by saturated models). *Let V and T be theories such that T is $PSC(\lambda)$ modulo V , for some $\lambda \geq \aleph_0$. Let \mathfrak{A} be a μ -saturated model of $V \cup T$, for some $\mu \geq \lambda$, and let \mathfrak{B} be a model of V . If \mathfrak{B} is sandwiched by \mathfrak{A} , then \mathfrak{B} is a model of T .*

Using the above lemmas, we can prove Theorem 5.2.1(1) as follows.

Proof of Theorem 5.2.1(1). We give the proof for theories without free variables. The proof for theories with free variables follows from definitions.

If: Suppose T is equivalent modulo V to a Σ_2^0 theory Y . We show that Y is $PSC(\lambda)$ modulo V for each $\lambda > \aleph_0$, whereby the same is true of T . Towards this, observe that Y is a countable set. Let \mathfrak{A} be a model of $V \cup Y$. Let $C \subseteq U_{\mathfrak{A}}$ be the (countable) set of witnesses in \mathfrak{A} , of the existential quantifiers of the sentences of Y . In other words, C is a countable subset of $U_{\mathfrak{A}}$ such that for each sentence ϕ of Y , there exist elements of C that form a witness in \mathfrak{A} , of the existential quantifiers of ϕ . It is easy to see that C is an \aleph_1 -crux of \mathfrak{A} w.r.t. Y modulo V . Then Y is $PSC(\aleph_1)$ modulo V , and hence $PSC(\lambda)$ modulo V , for each $\lambda > \aleph_0$.

Only If: Suppose T is $PSC(\lambda)$ modulo V for $\lambda > \aleph_0$. To complete the proof, it suffices to show, owing to Theorem 5.2.4, that T is preserved under sandwiches by models of T modulo V . Suppose \mathfrak{A}_1 and \mathfrak{B}_1 are given structures such that \mathfrak{B}_1 is sandwiched by \mathfrak{A}_1 , $\mathfrak{B}_1 \models V$ and $\mathfrak{A}_1 \models (V \cup T)$. Consider a μ -saturated elementary extension \mathfrak{A} of \mathfrak{A}_1 , for some $\mu \geq \lambda$. By Lemma 5.2.5, there exists a structure \mathfrak{B} isomorphic to \mathfrak{B}_1 such that \mathfrak{B} is sandwiched by \mathfrak{A} . Then $\mathfrak{A} \models (V \cup T)$ and $\mathfrak{B} \models V$, whence by Lemma 5.2.6, we have $\mathfrak{B} \models T$. Since $\mathfrak{B}_1 \cong \mathfrak{B}$, we have $\mathfrak{B}_1 \models T$, completing the proof. \square

We now prove Lemmas 5.2.5 and 5.2.6. We refer the reader to Section 2.3 for the notions of $\tau_{\mathfrak{A}}$, $\mathfrak{A}_{\mathfrak{A}}$, $\mathfrak{B}_{\mathfrak{A}}$, $\text{Diag}(\mathfrak{A})$ and $\text{El-diag}(\mathfrak{A})$ for $\mathfrak{A} \subseteq \mathfrak{B}$, that we use in our proofs below. We make the simple yet important observation that each of $\text{Diag}(\mathfrak{A})$ and $\text{El-diag}(\mathfrak{A})$ is closed under finite conjunctions. We let $\mathfrak{A} \preceq_1 \mathfrak{B}$ denote that (i) $\mathfrak{A} \subseteq \mathfrak{B}$ and (ii) every Σ_1^0 sentence of $\text{FO}(\tau_{\mathfrak{A}})$ that is true in $\mathfrak{B}_{\mathfrak{A}}$ is also true in $\mathfrak{A}_{\mathfrak{A}}$.

Lemma 5.2.7. $\mathfrak{A} \preceq_1 \mathfrak{B}$ iff there exists \mathfrak{A}' such that $(\mathfrak{A}, \mathfrak{B}, \mathfrak{A}')$ is a sandwich.

Proof. The ‘If’ direction follows easily from the definition of elementary substructure and the fact that existential formulae are preserved under extensions. For the converse, suppose that $\mathfrak{A} \preceq_1 \mathfrak{B}$. Let the vocabularies $\tau_{\mathfrak{B}}$ and $\tau_{\mathfrak{A}}$ be such that for every element a of \mathfrak{A} , the constant in $\tau_{\mathfrak{B}}$ corresponding to a is *the same* as the constant in $\tau_{\mathfrak{A}}$ corresponding to a (and hence the constants in $\tau_{\mathfrak{B}} \setminus \tau_{\mathfrak{A}}$ correspond exactly to the elements in \mathfrak{B} that are not in \mathfrak{A}). Now consider the theory Y given by $Y = \text{Diag}(\mathfrak{B}) \cup \text{El-diag}(\mathfrak{A})$. Any non-empty finite subset of $\text{Diag}(\mathfrak{B})$, resp. $\text{El-diag}(\mathfrak{A})$, is satisfied in $\mathfrak{B}_{\mathfrak{B}}$, resp. $\mathfrak{A}_{\mathfrak{A}}$. Let Z be any finite subset of Y , that has a non-empty intersection with both $\text{Diag}(\mathfrak{B})$ and $\text{El-diag}(\mathfrak{A})$; we can consider Z as given by $Z = \{\xi, \psi\}$ where $\xi \in \text{Diag}(\mathfrak{B})$ and $\psi \in \text{El-diag}(\mathfrak{A})$. Let c_1, \dots, c_r be the (distinct) constants of $\tau_{\mathfrak{B}} \setminus \tau_{\mathfrak{A}}$ appearing in ξ , and let x_1, \dots, x_r be fresh variables. Consider the sentence ϕ given by $\phi = \exists x_1 \dots \exists x_r \xi [c_1 \mapsto x_1; \dots; c_r \mapsto x_r]$, where $c_i \mapsto x_i$ denotes substitution of x_i for c_i , for $1 \leq i \leq r$. Observe that ϕ is a Σ_1^0 sentence of $\text{FO}(\tau_{\mathfrak{A}})$ and that $\mathfrak{B}_{\mathfrak{A}} \models \phi$. Since $\mathfrak{A} \preceq_1 \mathfrak{B}$, we have that $\mathfrak{A}_{\mathfrak{A}} \models \phi$. Let a_1, \dots, a_r be the witnesses in $\mathfrak{A}_{\mathfrak{A}}$, of the quantifiers of ϕ corresponding to variables x_1, \dots, x_r . Interpreting the constants c_1, \dots, c_r as a_1, \dots, a_r respectively, we see that $(\mathfrak{A}_{\mathfrak{A}}, a_1, \dots, a_r) \models Z$. Since Z is an arbitrary finite subset of Y , by the compactness theorem, Y is satisfied in a $\tau_{\mathfrak{B}}$ -structure \mathfrak{C} . The τ -reduct of \mathfrak{C} is the desired structure \mathfrak{A}' . \square

Proof of Lemma 5.2.5. Let \mathfrak{A} be a μ -saturated elementary extension of \mathfrak{A}_1 , for some $\mu \geq \omega$. We show below the existence of a structure \mathfrak{B}_2 such that (i) $\mathfrak{A} \preceq_1 \mathfrak{B}_2$ and (ii) \mathfrak{B}_1 is elementarily embeddable in \mathfrak{B}_2 via an embedding say f . Let \mathfrak{B} be the image of \mathfrak{B}_1 under f ; then $\mathfrak{B} \cong \mathfrak{B}_1$ and $\mathfrak{B} \preceq \mathfrak{B}_2$. By Lemma 5.2.7, there exists a structure \mathfrak{A}_2 such that $(\mathfrak{A}, \mathfrak{B}_2, \mathfrak{A}_2)$ is a sandwich, whence \mathfrak{B} is sandwiched by \mathfrak{A} . Then \mathfrak{B} is indeed as desired. For our arguments below, we make the following observation, call it (*): If \mathfrak{B} is sandwiched by \mathfrak{A} , then every Σ_2^0 sentence true in \mathfrak{A} is also true in \mathfrak{B} . This follows simply from Theorem 5.2.4 by taking T to be the set of all Σ_2^0 sentences that are true in \mathfrak{A} , and taking V to be the empty theory.

Let τ be the vocabulary of \mathfrak{A} and \mathfrak{B}_1 , and let $\tau_{\mathfrak{A}}$ and $\tau_{\mathfrak{B}_1}$ be such that $\tau_{\mathfrak{A}} \cap \tau_{\mathfrak{B}_1} = \tau$. Consider

the theory Y given by $Y = S_{\Pi}(\mathfrak{A}_{\mathfrak{A}}) \cup \text{El-diag}(\mathfrak{B}_1)$, where $S_{\Pi}(\mathfrak{A}_{\mathfrak{A}})$ denotes the set of all Π_1^0 sentences true in $\mathfrak{A}_{\mathfrak{A}}$. Observe that $S_{\Pi}(\mathfrak{A}_{\mathfrak{A}})$ is closed under finite conjunctions. Let Z be any non-empty finite subset of Y . If $Z \subseteq S_{\Pi}(\mathfrak{A}_{\mathfrak{A}})$ or $Z \subseteq \text{El-diag}(\mathfrak{B}_1)$, then Z is clearly satisfiable. Else, $Z = \{\xi, \psi\}$ where $\xi \in S_{\Pi}(\mathfrak{A}_{\mathfrak{A}})$ and $\psi \in \text{El-diag}(\mathfrak{B}_1)$. Let c_1, \dots, c_r be the (distinct) constants of $\tau_{\mathfrak{A}} \setminus \tau$ appearing in ξ , and let x_1, \dots, x_r be fresh variables. Consider the sentence ϕ given by $\phi = \exists x_1 \dots \exists x_r \xi [c_1 \mapsto x_1; \dots; c_r \mapsto x_r]$, where $c_i \mapsto x_i$ denotes substitution of x_i for c_i , for $1 \leq i \leq r$. Clearly $\mathfrak{A} \models \phi$, whence $\mathfrak{A}_1 \models \phi$. Since \mathfrak{B}_1 is sandwiched by \mathfrak{A}_1 and ϕ is a Σ_2^0 sentence, it follows from observation (*) above, that $\mathfrak{B}_1 \models \phi$. Let b_1, \dots, b_r be the witnesses in \mathfrak{B}_1 of the quantifiers of ϕ associated with x_1, \dots, x_r . One can now check that if $\mathfrak{R} = \mathfrak{B}_1$, then $(\mathfrak{R}, b_1, \dots, b_r) \models Z$. Since Z is an arbitrary finite subset of Y , by compactness theorem, Y is satisfiable. Whereby, there exists a τ -structure \mathfrak{B}_2 such that (i) $\mathfrak{A} \preceq_1 \mathfrak{B}_2$ and (ii) \mathfrak{B}_1 is elementarily embeddable in \mathfrak{B}_2 . \square

We now turn to proving Lemma 5.2.6. The notion of the FO-type of a k -tuple in a given structure for $k \in \mathbb{N}$ can be naturally extended to the notion of the FO-type of a tuple of length $< \lambda$ in a given structure for $\lambda \geq \aleph_0$. Formally, given a structure \mathfrak{A} and a tuple $\bar{a} = (a_1, a_2, \dots)$ of \mathfrak{A} , of length $< \lambda$, the FO-type of \bar{a} in \mathfrak{A} , denoted $\text{tp}_{\mathfrak{A}, \bar{a}}(x_1, x_2, \dots)$, is the set of formulae given by $\text{tp}_{\mathfrak{A}, \bar{a}}(x_1, x_2, \dots) = \{\varphi(x_{\eta_1}, \dots, x_{\eta_k}) \mid k \in \mathbb{N}, 1 \leq \eta_1 < \dots < \eta_k < \lambda, \varphi(x_{\eta_1}, \dots, x_{\eta_k}) \text{ is an FO formula such that } (\mathfrak{A}, a_{\eta_1}, \dots, a_{\eta_k}) \models \varphi(x_{\eta_1}, \dots, x_{\eta_k})\}$. The subset of $\text{tp}_{\mathfrak{A}, \bar{a}}(x_1, x_2, \dots)$ consisting of all Π_1^0 formulae in $\text{tp}_{\mathfrak{A}, \bar{a}}(x_1, x_2, \dots)$ is denoted as $\text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, x_2, \dots)$. For theories V and T such that T is $PSC(\lambda)$ modulo V , and for \mathfrak{A} and \bar{a} as mentioned above, we say that $\text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, x_2, \dots)$ *determines a λ -crux w.r.t. T modulo V* if it is the case that given a model \mathfrak{D} of V and a tuple \bar{d} of \mathfrak{D} , of length equal to that of \bar{a} , if $(\mathfrak{D}, \bar{d}) \models \text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, x_2, \dots)$, then $\mathfrak{D} \models T$. Since universal formulae are preserved under substructures, it follows that for \mathfrak{D} as just mentioned, the elements of \bar{d} form a λ -crux of \mathfrak{D} w.r.t. T modulo V . To prove Lemma 5.2.6, we need the next result which characterizes when a Π_1^0 -type determines a λ -crux.

Lemma 5.2.8 (Characterizing “crux determination”). *Let V and T be theories such that T is $PSC(\lambda)$ modulo V for some $\lambda \geq \aleph_0$. Let \mathfrak{A} be a model of V and \bar{a} be a tuple of elements of \mathfrak{A} , of length less than λ . Then $\text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, x_2, \dots)$ determines a λ -crux w.r.t. T modulo V iff $\mathfrak{A} \models T$ and for some $\mu \geq \lambda$, there exists a μ -saturated elementary extension \mathfrak{B} of \mathfrak{A} (hence $\mathfrak{B} \models (V \cup T)$) such that \bar{a} is a λ -crux of \mathfrak{B} w.r.t. T modulo V .*

Proof. ‘Only If:’ Since $\mathfrak{A} \models V$ and $(\mathfrak{A}, \bar{a}) \models \text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, x_2, \dots)$, we have that $\mathfrak{A} \models T$. Let \mathfrak{B} be any μ -saturated elementary extension of \mathfrak{A} for $\mu \geq \lambda$; then $(\mathfrak{B}, \bar{a}) \models \text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, x_2, \dots)$ and $\mathfrak{B} \models V$. Since $\text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, x_2, \dots)$ determines a λ -crux w.r.t. T modulo V , we have that \bar{a} is a λ -crux of \mathfrak{B} w.r.t. T modulo V .

‘If:’ Let \mathfrak{A} , \mathfrak{B} and \bar{a} be as mentioned in the statement. Consider a model \mathfrak{D} of V and a tuple \bar{d} of \mathfrak{D} , of length equal to that of \bar{a} , such that $(\mathfrak{D}, \bar{d}) \models \text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, x_2, \dots)$. By the downward Löwenheim-Skolem theorem, there exists $\mathfrak{D}_1 \preceq \mathfrak{D}$ such that (i) \mathfrak{D}_1 contains \bar{d} and (ii) $|\mathfrak{D}_1| \leq \lambda$. Then $(\mathfrak{D}_1, \bar{d}) \models \text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, x_2, \dots)$. Now since $\mathfrak{A} \preceq \mathfrak{B}$, we have that $\text{tp}_{\Pi, \mathfrak{B}, \bar{a}}(x_1, x_2, \dots) = \text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, x_2, \dots)$. Then every existential sentence that is true in $(\mathfrak{D}_1, \bar{d})$ is also true in (\mathfrak{B}, \bar{a}) . Since \mathfrak{B} is μ -saturated, and the length of \bar{a} is $< \lambda \leq \mu$, we have that (\mathfrak{B}, \bar{a}) is also μ -saturated (by Proposition 2.4.2(1)). Further, since $|\mathfrak{D}_1| \leq \lambda$, we have $|(\mathfrak{D}_1, \bar{d})| \leq \lambda \leq \mu$. Then there exists an embedding $f : (\mathfrak{D}_1, \bar{d}) \rightarrow (\mathfrak{B}, \bar{a})$ (by Proposition 2.4.2(5)). The image of $(\mathfrak{D}_1, \bar{d})$ under f is a substructure $(\mathfrak{B}_1, \bar{a})$ of (\mathfrak{B}, \bar{a}) . Since $\mathfrak{D}_1 \preceq \mathfrak{D}$ and $\mathfrak{D} \models V$, we have $\mathfrak{B}_1 \models V$. Further since \bar{a} forms a λ -crux of \mathfrak{B} w.r.t. T modulo V (by assumption), we have $\mathfrak{B}_1 \models T$. Then \mathfrak{D}_1 , and hence \mathfrak{D} , models T , completing the proof. \square

Proof of Lemma 5.2.6. We assume the vocabulary to be τ . Since \mathfrak{B} is sandwiched by \mathfrak{A} , there exist structures \mathfrak{A}_1 and \mathfrak{B}_1 such that (i) $\mathfrak{B} \preceq \mathfrak{B}_1$ and (ii) $(\mathfrak{A}, \mathfrak{B}_1, \mathfrak{A}_1)$ is a sandwich. Let \mathfrak{D} be a μ -saturated elementary extension of \mathfrak{A}_1 for some $\mu \geq \lambda$. Then $\mathfrak{A} \preceq \mathfrak{D}$. Since \mathfrak{A} models $V \cup T$, so does \mathfrak{D} .

Now, given that T is $PSC(\lambda)$ modulo V , there exists a λ -crux of \mathfrak{D} w.r.t. T modulo V ; let \bar{d} be any tuple (of length $< \lambda$) formed from this λ -crux. Consider $\text{tp}_{\mathfrak{D}, \bar{d}}(x_1, x_2, \dots)$, namely the FO-type of \bar{d} in \mathfrak{D} . Since $\mathfrak{A} \preceq \mathfrak{D}$, we have $\mathfrak{A} \equiv \mathfrak{D}$ (see Lemma 2.3.1). Then since \mathfrak{A} is μ -saturated, there exists a tuple \bar{a} of \mathfrak{A} , of length equal to that of \bar{d} , such that $(\mathfrak{A}, \bar{a}) \equiv (\mathfrak{D}, \bar{d})$ (by Proposition 2.4.2(4)). In other words, $\text{tp}_{\mathfrak{A}, \bar{a}}(x_1, x_2, \dots) = \text{tp}_{\mathfrak{D}, \bar{d}}(x_1, x_2, \dots)$. Then since $\mathfrak{A} \preceq \mathfrak{D}$, it follows that the FO-type of \bar{a} in \mathfrak{D} , namely $\text{tp}_{\mathfrak{D}, \bar{a}}(x_1, x_2, \dots)$, is exactly $\text{tp}_{\mathfrak{A}, \bar{a}}(x_1, x_2, \dots)$. Whence, $\text{tp}_{\Pi, \mathfrak{D}, \bar{a}}(x_1, x_2, \dots) = \text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, x_2, \dots) = \text{tp}_{\Pi, \mathfrak{D}, \bar{d}}(x_1, x_2, \dots)$. Now since (i) $\mathfrak{D} \models (V \cup T)$, (ii) \mathfrak{D} is itself μ -saturated and (iii) \bar{d} is a λ -crux of \mathfrak{D} w.r.t. T modulo V , we have by Lemma 5.2.8, that $\text{tp}_{\Pi, \mathfrak{D}, \bar{d}}(x_1, x_2, \dots)$, and hence $\text{tp}_{\Pi, \mathfrak{D}, \bar{a}}(x_1, x_2, \dots)$, determines a λ -crux w.r.t. T modulo V . Whence the elements of \bar{a} form a λ -crux of \mathfrak{D} w.r.t. T modulo V . Since (i) $\mathfrak{B}_1 \subseteq \mathfrak{A}_1 \preceq \mathfrak{D}$ (ii) \mathfrak{B}_1 contains \bar{a} and (iii) $\mathfrak{B}_1 \models V$ (since $\mathfrak{B} \models V$ and $\mathfrak{B} \preceq \mathfrak{B}_1$), we have by definition of a λ -crux (w.r.t. T modulo V), that $\mathfrak{B}_1 \models T$, whence $\mathfrak{B} \models T$. \square

The characterizations of this section are depicted pictorially below.

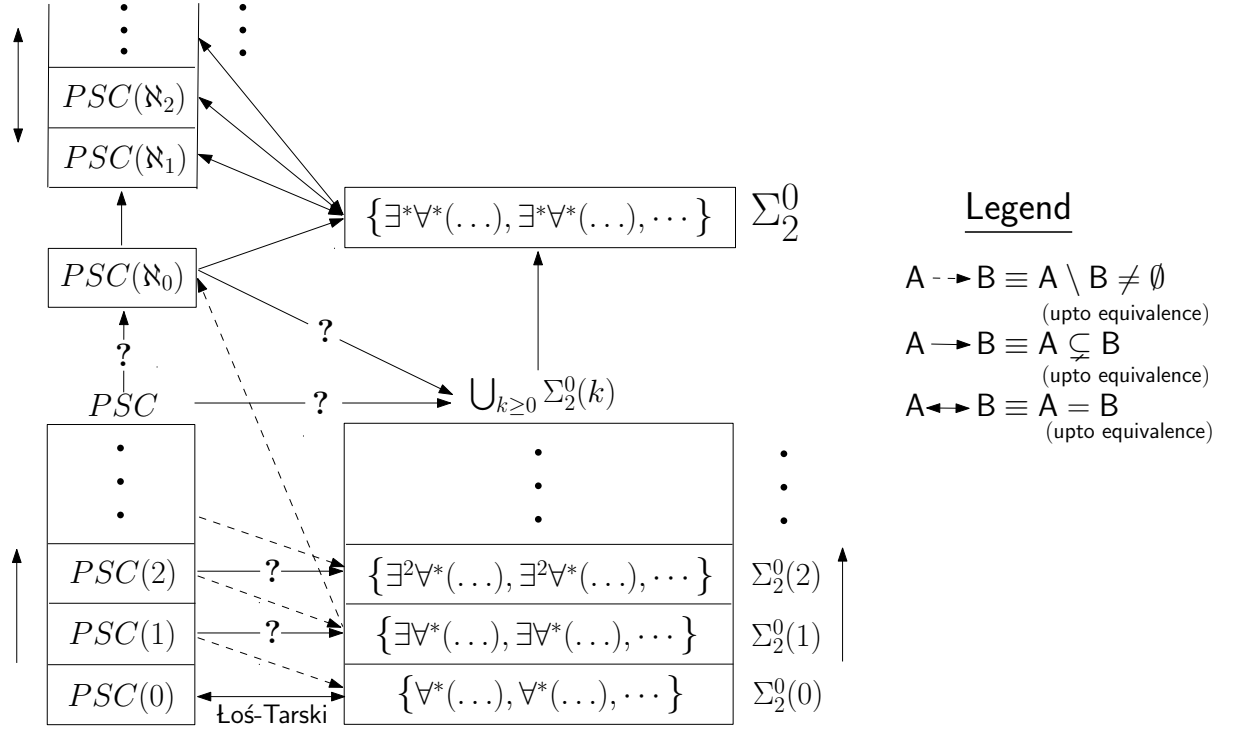


Figure 5.2: (Partial) characterizations of $PSC(k)$ and $PSC(\lambda)$ theories

Proof of Theorem 5.2.3

The technique of our proof is as presented below.

1. We first define a variant of $PSC(k)$, that we call $PSC_{var}(k)$, into whose definition we build Hypothesis 5.2.2.
2. We then show that $PSC_{var}(k)$ theories are equivalent to theories of $\exists^k \forall^*$ sentences. This is done in the following two steps:
 - “Going up”: We give a characterization of $PSC_{var}(k)$ theories in terms of sentences of a special infinitary logic (Lemma 5.2.15).
 - “Coming down”: We provide a translation of sentences of the aforesaid infinitary logic, into their equivalent FO theories, whenever these sentences define *elementary* (i.e. definable using FO theories) classes of structures (Proposition 5.2.16). The FO theories are obtained from suitable *finite approximations* of the infinitary sentences, and turn out to be theories of $\exists^k \forall^*$ sentences.
3. We hypothesize that $PSC_{var}(k)$ theories are no different from $PSC(k)$ theories, as an equivalent reformulation of Hypothesis 5.2.2, to obtain Theorem 5.2.3. To show that this hypoth-

esis is well-motivated, we define a variant of $PCE(k)$, denoted $PCE_{var}(k)$, that is dual to $PSC_{var}(k)$. We show that $PCE_{var}(k)$ coincides with $PCE(k)$ for theories, and use this to conclude that $PSC_{var}(k)$ coincides with $PSC(k)$ for sentences (Lemma 5.2.12).

Throughout the section, whenever V and T are clear from the context, we skip mentioning the qualifier ‘w.r.t. T modulo V ’ for a k -crux, if T is $PSC(k)$ modulo V . Before we present the definitions of $PSC_{var}(k)$ and $PCE_{var}(k)$, we first define the notion of ‘distinguished k -crux’.

Definition 5.2.9. Suppose T is $PSC(k)$ modulo V for theories T and V . Given a model \mathfrak{A} of $V \cup T$, we call a k -tuple \bar{a} of \mathfrak{A} a *distinguished k -crux* of \mathfrak{A} , if for some $\mu \geq \omega$, there is a μ -saturated elementary extension \mathfrak{A}^+ of \mathfrak{A} (whence $\mathfrak{A}^+ \models V \cup T$) such that \bar{a} is a k -crux of \mathfrak{A}^+ (whence \bar{a} is also a k -crux of \mathfrak{A}).

We now define $PSC_{var}(k)$ and $PCE_{var}(k)$. We refer the reader to Definition 4.1.4 for the meaning of the phrase ‘ k -ary cover of \mathfrak{A} in \mathfrak{A}^+ ’ appearing in the definition below.

Definition 5.2.10. Let V and T be theories.

1. We say T is $PSC_{var}(k)$ modulo V if T is $PSC(k)$ modulo V and every model of $V \cup T$ contains a distinguished k -crux.
2. We say T is $PCE_{var}(k)$ modulo V if for every model \mathfrak{A} of V , there exists a μ -saturated elementary extension \mathfrak{A}^+ of \mathfrak{A} for some $\mu \geq \omega$, such that for every collection R of models of $V \cup T$, if R is a k -ary cover of \mathfrak{A} in \mathfrak{A}^+ , then $\mathfrak{A} \models T$.

If $\phi(\bar{x})$ and $T(\bar{x})$ are respectively a formula and a theory, each of whose free variables are among \bar{x} , then for a theory V , the notions of ‘ $\phi(\bar{x})$ is $PSC_{var}(k)$ (resp. $PCE_{var}(k)$) modulo V ’ and ‘ $T(\bar{x})$ is $PSC_{var}(k)$ (resp. $PCE_{var}(k)$) modulo V ’ are defined similar to corresponding notions for $PSC(k)$ (resp. $PCE(k)$). The following duality is easy to see.

Lemma 5.2.11 ($PSC_{var}(k)$ - $PCE_{var}(k)$ duality). *Given a theory V , a formula $\phi(\bar{x})$ is $PSC_{var}(k)$ modulo V iff $\neg\phi(\bar{x})$ is $PCE_{var}(k)$ modulo V .*

Towards the proof of Theorem 5.2.3, we first show the following.

Lemma 5.2.12. *Given a theory V , each of the following holds.*

1. A formula $\phi(\bar{x})$ is $PSC(k)$ modulo V iff $\phi(\bar{x})$ is $PSC_{var}(k)$ modulo V .
2. A theory $T(\bar{x})$ is $PCE(k)$ modulo V iff $T(\bar{x})$ is $PCE_{var}(k)$ modulo V .

Proof. We show below the following equivalence, call it (\dagger) : A theory $T(\bar{x})$ is $PCE_{var}(k)$ modulo V iff $T(\bar{x})$ is equivalent modulo V to a theory of $\forall^k\exists^*$ formulae, all of whose free variables are among \bar{x} . Then part (2) of this lemma follows from (\dagger) and Theorem 5.1.1(1). Part (1) of the lemma in turn follows from part (2) and the dualities given by Lemma 3.2.5 and Lemma 5.2.11.

Given that the notion of ‘ k -ary cover of \mathfrak{A} in \mathfrak{A}' ’ is the same as the notion of ‘ k -ary cover’ as defined in Definition 3.2.1, we can prove the ‘Only if’ direction of (\dagger) in a manner identical to the proof of the ‘Only if’ direction of Theorem 5.1.1(1). The proof of the ‘If’ direction of (\dagger) is also nearly the same as that of the ‘If’ direction of Theorem 5.1.1(1); we present this proof below for completeness. It suffices to give the proof for theories without free variables.

Let T be equivalent modulo V to a theory of $\forall^k\exists^*$ sentences. Given a model \mathfrak{A} of V , let \mathfrak{A}^+ be a μ -saturated elementary extension of \mathfrak{A} , for some $\mu \geq \omega$. Let R be a collection of models of $V \cup T$ that forms a k -ary cover of \mathfrak{A} in \mathfrak{A}^+ . We show that $\mathfrak{A} \models T$. Consider $\varphi \in T$; let $\varphi = \forall^k \bar{x} \psi(\bar{x})$ for a Σ_1^0 formula $\psi(\bar{x})$, and let \bar{a} be a k -tuple of \mathfrak{A} . Since R is a k -ary cover of \mathfrak{A} in \mathfrak{A}^+ , there exists $\mathfrak{B}_{\bar{a}} \in R$ such that $\mathfrak{B}_{\bar{a}}$ contains \bar{a} . Since $\mathfrak{B}_{\bar{a}} \models (V \cup T)$, we have $\mathfrak{B}_{\bar{a}} \models \varphi$ and hence $(\mathfrak{B}_{\bar{a}}, \bar{a}) \models \psi(\bar{x})$. Since $\psi(\bar{x})$ is a Σ_1^0 formula and $\mathfrak{B}_{\bar{a}} \subseteq \mathfrak{A}^+$, we have $(\mathfrak{A}^+, \bar{a}) \models \psi(\bar{x})$, whence $(\mathfrak{A}, \bar{a}) \models \psi(\bar{x})$ since $\mathfrak{A} \preceq \mathfrak{A}^+$. Since \bar{a} is arbitrary, $\mathfrak{A} \models \varphi$, and since φ is an arbitrary sentence of T , we have $\mathfrak{A} \models T$. \square

Motivated by Lemma 5.2.12, we put forth the hypothesis below.

Hypothesis 5.2.13. *If V and $T(\bar{x})$ are theories, then $T(\bar{x})$ is $PSC(k)$ modulo V iff $T(\bar{x})$ is $PSC_{var}(k)$ modulo V .*

It is easy to see that Hypothesis 5.2.13 is an equivalent reformulation of Hypothesis 5.2.2.

The following result is the essence of Theorem 5.2.3.

Theorem 5.2.14. *Given theories V and $T(\bar{x})$, suppose $T(\bar{x})$ is $PSC_{var}(k)$ modulo V . Then $T(\bar{x})$ is equivalent modulo V to a theory of Σ_2^0 formulae, all of whose free variables are among \bar{x} , and all of which have k existential quantifiers.*

Proof of Theorem 5.2.3. Follows from the equivalence of Hypotheses 5.2.2 and 5.2.13, and Theorem 5.2.14. \square

We devote the rest of this section to proving Theorem 5.2.14. We first introduce some notation and terminology. These are adapted versions of similar notation and terminology introduced

in [45] and [46]. Given a class \mathcal{F} of formulae and $k \geq 0$, denote by $[\exists^k \wedge] \mathcal{F}$ the class of infinitary formulae $\Phi(\bar{x})$ of the form $\exists y_1 \dots \exists y_k \bigwedge_{i \in I} \psi_i(y_1, \dots, y_k, \bar{x})$ where I is an index set (of arbitrary cardinality) and for each $i \in I$, ψ_i is a formula of \mathcal{F} , whose free variables are among y_1, \dots, y_k, \bar{x} . Let $[\exists^* \wedge] \mathcal{F} = \bigcup_{k \geq 0} [\exists^k \wedge] \mathcal{F}$. Observe that $\mathcal{F} \subseteq [\exists^* \wedge] \mathcal{F}$. For each $j \in \mathbb{N}$, let $[\exists^* \wedge]^j \mathcal{F} = [\exists^* \wedge] [\exists^* \wedge]^{j-1} \mathcal{F}$, where $[\exists^* \wedge]^0 \mathcal{F} = \mathcal{F}$. Let $[\exists^* \wedge]^* \mathcal{F} = \bigcup_{j \geq 0} [\exists^* \wedge]^j \mathcal{F}$. Finally, let $[\vee] \mathcal{F}$ denote arbitrary disjunctions of formulae of \mathcal{F} . Observe that $\mathcal{F} \subseteq [\vee] \mathcal{F}$. Let $\Phi(\bar{x})$ be a formula of $[\vee] [\exists^* \wedge]^* \text{FO}$, where FO denotes as usual, the class of all first order formulae. We define below, the set $\mathcal{A}(\Phi)(\bar{x})$ of *finite approximations* of $\Phi(\bar{x})$. Let \subseteq_f denote ‘finite subset of’.

1. If $\Phi(\bar{x}) \in \text{FO}$, then $\mathcal{A}(\Phi)(\bar{x}) = \{\Phi(\bar{x})\}$.
2. If $\Phi(\bar{x}) = \exists^k \bar{y} \bigwedge_{i \in I} \Psi_i(\bar{x}, \bar{y})$ for $k \geq 0$ and some index set I , then $\mathcal{A}(\Phi)(\bar{x}) = \{\exists^k \bar{y} \bigwedge_{i \in I_1} \gamma_i(\bar{x}, \bar{y}) \mid \gamma_i(\bar{x}, \bar{y}) \in \mathcal{A}(\Psi_i)(\bar{x}, \bar{y}), I_1 \subseteq_f I\}$.
3. If $\Phi(\bar{x}) = \bigvee_{i \in I} \Psi_i(\bar{x})$, then $\mathcal{A}(\Phi)(\bar{x}) = \{\bigvee_{i \in I_1} \gamma_i(\bar{x}) \mid \gamma_i(\bar{x}) \in \mathcal{A}(\Psi_i)(\bar{x}), I_1 \subseteq_f I\}$.

Our proof of Theorem 5.2.14 is in two parts. The first part, namely the “going up” part as alluded to in the beginning of this subsection, gives a characterization of $PSC_{var}(k)$ theories in terms of the formulae of $[\vee] [\exists^k \wedge] \Pi_1^0$, where Π_1^0 is the usual class of all prenex FO formulae having only universal quantifiers.

Lemma 5.2.15. *Let V and $T(\bar{x})$ be given theories. Then $T(\bar{x})$ is $PSC_{var}(k)$ modulo V iff $T(\bar{x})$ is equivalent modulo V to a formula of $[\vee] [\exists^k \wedge] \Pi_1^0$, whose free variables are among \bar{x} .*

The second part of the proof of Theorem 5.2.14, namely the “coming down” part, consists of getting FO theories equivalent to the formulae of $[\vee] [\exists^k \wedge] \Pi_1^0$, whenever the latter define elementary classes of structures. In fact, we show a more general result as we now describe. Given a theory V , we say that a formula $\Phi(x_1, \dots, x_k)$ of $[\vee] [\exists^* \wedge]^* \text{FO}$ (over a vocabulary say τ) *defines an elementary class modulo V* if the sentence (over the vocabulary τ_k) obtained by substituting fresh and distinct constants c_1, \dots, c_k for the free occurrences of x_1, \dots, x_k in $\Phi(x_1, \dots, x_k)$, defines an elementary class (of τ_k -structures) modulo V . The result below characterizes formulae of $[\vee] [\exists^* \wedge]^* \text{FO}$ that define elementary classes, in terms of the finite approximations of these formulae.

Proposition 5.2.16. *Let $\Phi(\bar{x})$ be a formula of $[\vee] [\exists^* \wedge]^* \text{FO}$ and V be a given theory. Then $\Phi(\bar{x})$ defines an elementary class modulo V iff $\Phi(\bar{x})$ is equivalent modulo V to a countable subset of $\mathcal{A}(\Phi)(\bar{x})$.*

The above results prove Theorem 5.2.14 as follows.

Proof of Theorem 5.2.14. For any formula $\Phi(\bar{x})$ of $[\bigvee] [\exists^k \bigwedge] \Pi_1^0$, each formula of the set $\mathcal{A}(\Phi)(\bar{x})$ can be seen to be equivalent to an $\exists^k \forall^*$ formula whose free variables are among \bar{x} . The result then follows from Lemma 5.2.15 and Proposition 5.2.16. \square

We now prove Lemma 5.2.15 and Proposition 5.2.16. We observe that it suffices to prove these results only for theories/formulae without free variables.

The proof of Lemma 5.2.15 requires the following result that characterizes when a k -crux of a model of a $PSC(k)$ theory is a distinguished k -crux of the model. To state this result, we define the notion of ‘the Π_1^0 -type of a k -tuple determining a k -crux’, analogously to the notion of the Π_1^0 -type of a tuple of length $< \lambda$ determining a λ -crux, that was introduced just before Lemma 5.2.8. Formally, given theories V and T such that T is $PSC(k)$ modulo V , and given a structure \mathfrak{A} and a k -tuple \bar{a} of \mathfrak{A} , we say $\text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(\bar{x})$ *determines a k -crux w.r.t. T modulo V* if it is the case that given a model \mathfrak{D} of V and a k -tuple \bar{d} of \mathfrak{D} , if $(\mathfrak{D}, \bar{d}) \models \text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(\bar{x})$, then $\mathfrak{D} \models T$. Since universal formulae are preserved under substructures, it follows that for \mathfrak{D} as just mentioned, the elements of \bar{d} form a k -crux of \mathfrak{D} w.r.t. T modulo V .

Lemma 5.2.17 (Characterizing distinguished k -cruxes). *Let V and T be theories such that T is $PSC(k)$ modulo V . Let \mathfrak{A} be a model of $V \cup T$ and \bar{a} be a k -tuple of elements of \mathfrak{A} . Then \bar{a} is a distinguished k -crux of \mathfrak{A} w.r.t. T modulo V iff $\text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(\bar{x})$ determines a k -crux w.r.t. T modulo V .*

Proof. Similar to the proof of Lemma 5.2.8 \square

Proof of Lemma 5.2.15. If: Let T be equivalent modulo V to the sentence $\Phi = \bigvee_{i \in I} \exists^k \bar{y}_i \bigwedge Y_i(\bar{y}_i)$, where I is an index set and for each $i \in I$, Y_i is a set of Π_1^0 formulae, all of whose free variables are among \bar{y}_i . Then given a model \mathfrak{A} of $V \cup T$, there exist $i \in I$ and \bar{a} in \mathfrak{A} such that $(\mathfrak{A}, \bar{a}) \models \bigwedge Y_i(\bar{y}_i)$. Let \mathfrak{A}^+ be a μ -saturated elementary extension of \mathfrak{A} , for some $\mu \geq \omega$. Then $(\mathfrak{A}^+, \bar{a}) \models \bigwedge Y_i(\bar{y}_i)$. Whence for each $\mathfrak{B} \subseteq \mathfrak{A}^+$ such that \mathfrak{B} contains \bar{a} , $(\mathfrak{B}, \bar{a}) \models \bigwedge Y_i(\bar{y}_i)$, and hence $\mathfrak{B} \models \Phi$. Since Φ is equivalent to T modulo V , we have \bar{a} as a distinguished k -crux of \mathfrak{A} . Only If: Suppose T is $PSC_{var}(k)$ modulo V . Given a model \mathfrak{A} of $V \cup T$, let $\text{Dist-}k\text{-cruxes}(\mathfrak{A})$ be the (non-empty) set of all distinguished k -cruxes of \mathfrak{A} . Consider the sentence $\Phi = \bigvee_{\mathfrak{A} \models V \cup T, \bar{a} \in \text{Dist-}k\text{-cruxes}(\mathfrak{A})} \exists^k \bar{x} \bigwedge \text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(\bar{x})$. We show that T is equivalent to Φ modulo

V . That T implies Φ modulo V is obvious from the definition of Φ . Towards the converse, suppose $\mathfrak{B} \models \{\Phi\} \cup V$. Then for some model \mathfrak{A} of $V \cup T$, some distinguished k -crux \bar{a} of \mathfrak{A} , and for some k -tuple \bar{b} of \mathfrak{B} , we have $(\mathfrak{B}, \bar{b}) \models \text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(\bar{x})$. By Lemma 5.2.17, $\text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(\bar{x})$ determines a k -crux w.r.t. T modulo V , whence $\mathfrak{B} \models T$. \square

We now turn to proving Proposition 5.2.16. Our proof of Proposition 5.2.16 crucially uses our compactness result for formulae of $[\exists^* \wedge]^* \text{FO}$, that we state now.

Lemma 5.2.18. *Let $\Phi(\bar{x})$ be a formula of $[\exists^* \wedge]^* \text{FO}$. If every formula of $\mathcal{A}(\Phi)(\bar{x})$ is satisfiable modulo a theory V , then $\Phi(\bar{x})$ is satisfiable modulo V .*

Observe that the standard compactness theorem for FO is a special case of the above result: Given an FO theory $T(\bar{x})$, let $\Phi(\bar{x})$ be the formula of $[\exists^* \wedge]^* \text{FO}$ given by $\Phi(\bar{x}) = \bigwedge T(\bar{x})$. Then every formula of $\mathcal{A}(\Phi)(\bar{x})$ is equivalent to a finite subset of $T(\bar{x})$ and vice-versa.

Remark 5.2.19. The formulas of $[\exists^* \wedge]^* \text{FO}$ are special kinds of “conjunctive formulas”, where the latter are as defined in [46]. The paper [46] gives a generalization of the compactness theorem by proving a compactness result for conjunctive formulas, whose statement is similar to that of Lemma 5.2.18. However, Lemma 5.2.18 does not follow from this result of [46] because the set of finite approximations of sentences $\Phi(\bar{x})$ of $[\exists^* \wedge]^* \text{FO}$, as defined in [46], is semantically strictly larger than the set $\mathcal{A}(\Phi)(\bar{x})$ that we have defined. Further, the techniques that we use in proving Lemma 5.2.18 are much different from those used in [46] for proving the compactness result for conjunctive formulas.

We finally require the following two auxiliary lemmas for the proofs of Proposition 5.2.16 and Lemma 5.2.18.

Lemma 5.2.20. *For $j \in \mathbb{N}$, let $T(\bar{x})$ be a set of formulae of $[\exists^* \wedge]^j \text{FO}$, all of whose free variables are among \bar{x} . If every finite subset of $T(\bar{x})$ is satisfiable modulo a theory V , then $T(\bar{x})$ is satisfiable modulo V .*

Proof. We prove the statement by induction on j . The base case of $j = 0$ is the standard compactness theorem. As induction hypothesis, suppose the statement is true for j . For the inductive step, consider a set $T(\bar{x}) = \{\Phi_i(\bar{x}) \mid i \in I\}$ of $[\exists^* \wedge]^{j+1} \text{FO}$ formulae, all of whose free variables are among \bar{x} , and suppose every finite subset of $T(\bar{x})$ is satisfiable modulo V . Let $\Phi_i(\bar{x}) = \exists \bar{y}_i \wedge T_i(\bar{x}, \bar{y}_i)$ where $T_i(\bar{x}, \bar{y}_i)$ is a set of formulae of $[\exists^* \wedge]^j \text{FO}$. Assume for $i, j \in I$

and $i \neq j$, that \bar{y}_i and \bar{y}_j have no common variables. We show that the set Y of $[\exists^* \wedge]^j$ FO formulae given by $Y = \bigcup_{i \in I} T_i$ is satisfiable modulo V ; then so is $T(\bar{x})$.

By the induction hypothesis, it suffices to show that every finite subset Z of Y is satisfiable modulo V . Let $Z(\bar{x}, \bar{y}_{i_1}, \dots, \bar{y}_{i_n}) = \bigcup_{r=1}^{r=n} Z_r(\bar{x}, \bar{y}_{i_r})$, where $n > 0$, $Z_r(\bar{x}, \bar{y}_{i_r}) \subseteq_f T_{i_r}(\bar{x}, \bar{y}_{i_r})$ and $i_r \in I$, for each $r \in \{1, \dots, n\}$. The subset $\{\Phi_{i_r}(\bar{x}) \mid r \in \{1, \dots, n\}\}$ of $T(\bar{x})$ is satisfiable modulo V by assumption, whence for some model \mathfrak{A} of V , and interpretations \bar{a} of \bar{x} and \bar{b}_{i_r} of \bar{y}_{i_r} , we have that $\bigcup_{r=1}^{r=n} T_{i_r}(\bar{x}, \bar{y}_{i_r})$ is satisfied in $(\mathfrak{A}, \bar{a}, \bar{b}_{i_1}, \dots, \bar{b}_{i_n})$; then $(\mathfrak{A}, \bar{a}, \bar{b}_{i_1}, \dots, \bar{b}_{i_n}) \models Z(\bar{x}, \bar{y}_{i_1}, \dots, \bar{y}_{i_n})$. \square

Lemma 5.2.21. *Let $\Phi(\bar{x})$ be a formula of $[\exists^* \wedge]^* \text{FO}$. If $(\mathfrak{A}, \bar{a}) \models \Phi(\bar{x})$, then $(\mathfrak{A}, \bar{a}) \models \xi(\bar{x})$ for every formula $\xi(\bar{x})$ of $\mathcal{A}(\Phi)(\bar{x})$.*

Proof. We prove the lemma by induction. The statement is trivial for formulae of $\text{FO} = [\exists^* \wedge]^0 \text{FO}$. Assume the statement for $[\exists^* \wedge]^j \text{FO}$ formulae. Consider a formula $\Phi(\bar{x})$ of $[\exists^* \wedge]^{j+1} \text{FO}$ given by $\Phi(\bar{x}) = \exists^n \bar{y} \bigwedge_{i \in I} \Psi_i(\bar{x}, \bar{y})$, where $\Psi_i(\bar{x}, \bar{y}) \in [\exists^* \wedge]^j \text{FO}$ for each $i \in I$. Consider a formula $\xi(\bar{x})$ of $\mathcal{A}(\Phi)(\bar{x})$; then $\xi(\bar{x}) = \exists^n \bar{y} \bigwedge_{i \in I_1} \gamma_i(\bar{x}, \bar{y})$, for some $I_1 \subseteq_f I$ and $\gamma_i(\bar{x}, \bar{y}) \in \mathcal{A}(\Psi_i)(\bar{x}, \bar{y})$ for each $i \in I_1$. Since $(\mathfrak{A}, \bar{a}) \models \Phi(\bar{x})$, there is an n -tuple \bar{b} from \mathfrak{A} such that $(\mathfrak{A}, \bar{a}, \bar{b}) \models \Psi_i(\bar{x}, \bar{y})$ for each $i \in I_1$. By induction hypothesis, $(\mathfrak{A}, \bar{a}, \bar{b}) \models \gamma_i(\bar{x}, \bar{y})$ for each $i \in I_1$; then $(\mathfrak{A}, \bar{a}) \models \xi(\bar{x})$. \square

Proof of Lemma 5.2.18. The proof proceeds by induction. The statement trivially holds for formulae of $\text{FO} = [\exists^* \wedge]^0 \text{FO}$. Assume the statement is true for formulae of $[\exists^* \wedge]^j \text{FO}$. Consider a formula $\Phi(\bar{x})$ of $[\exists^* \wedge]^{j+1} \text{FO}$ given by $\Phi(\bar{x}) = \exists \bar{y} \bigwedge_{i \in I} \Psi_i(\bar{x}, \bar{y})$, where $\Psi_i(\bar{x}, \bar{y})$ is a formula of $[\exists^* \wedge]^j \text{FO}$ for each $i \in I$. We show that every finite subset of $T(\bar{x}, \bar{y}) = \{\Psi_i(\bar{x}, \bar{y}) \mid i \in I\}$ is satisfiable modulo V . Then by Lemma 5.2.20, $T(\bar{x}, \bar{y})$ is satisfiable modulo V ; then $\Phi(\bar{x})$ is also satisfiable modulo V .

Let I_1 be a finite subset of I . For $i \in I_1$, consider the formula $\Psi_i(\bar{x}, \bar{y})$ of $T(\bar{x}, \bar{y})$; it is given by $\Psi_i(\bar{x}, \bar{y}) = \exists \bar{z}_i \bigwedge Z_i(\bar{x}, \bar{y}, \bar{z}_i)$ where $Z_i(\bar{x}, \bar{y}, \bar{z}_i)$ is a set of formulas of $[\exists^* \wedge]^{j-1} \text{FO}$. Let $\bar{z} = (\bar{z}_i)_{i \in I_1}$ be the tuple of all the variables of the \bar{z}_i s, for i ranging over I_1 . Assume without loss of generality that for $i_1, i_2 \in I$ such that $i_1 \neq i_2$, none of the variables of \bar{z}_{i_1} appear in Ψ_{i_2} . Consider the formula $\Psi(\bar{x}, \bar{y})$ of $[\exists^* \wedge]^j \text{FO}$ given by $\Psi(\bar{x}, \bar{y}) = \exists \bar{z} \bigwedge (\bigcup_{i \in I_1} Z_i(\bar{x}, \bar{y}, \bar{z}_i))$. It is easy to verify that $\Psi(\bar{x}, \bar{y})$ is equivalent (over all structures) to $\{\Psi_i(\bar{x}, \bar{y}) \mid i \in I_1\}$. We now show that the latter is satisfiable modulo V by showing that the former is satisfiable modulo V – this in turn is done by showing that every formula in $\mathcal{A}(\Psi)(\bar{x}, \bar{y})$ is satisfiable modulo V , and

then applying the induction hypothesis mentioned at the outset.

Let $\gamma(\bar{x}, \bar{y})$ be an arbitrary formula of $\mathcal{A}(\Psi)(\bar{x}, \bar{y})$. Then $\gamma(\bar{x}, \bar{y})$ is of the form $\exists \bar{z} \bigwedge_{i \in I_2} \bigwedge_{l \in \{1, \dots, n_i\}} \alpha_{i,l}(\bar{x}, \bar{y}, \bar{z}_i)$, where $I_2 \subseteq I_1$, and for each $i \in I_2$, we have $n_i \geq 1$, $\alpha_{i,l}(\bar{x}, \bar{y}, \bar{z}_i) \in \mathcal{A}(\beta_{i,l})(\bar{x}, \bar{y}, \bar{z}_i)$, and $\{\beta_{i,1}(\bar{x}, \bar{y}, \bar{z}_i), \dots, \beta_{i,n_i}(\bar{x}, \bar{y}, \bar{z}_i)\} \subseteq_f Z_i(\bar{x}, \bar{y}, \bar{z}_i)$. It is easy to see that $\gamma(\bar{x}, \bar{y})$ is equivalent to the formula $\bigwedge_{i \in I_2} \gamma_i(\bar{x}, \bar{y})$ where $\gamma_i(\bar{x}, \bar{y}) = \exists \bar{z}_i \bigwedge_{l \in \{1, \dots, n_i\}} \alpha_{i,l}(\bar{x}, \bar{y}, \bar{z}_i)$. Observe now that $\gamma_i(\bar{x}, \bar{y}) \in \mathcal{A}(\Psi_i)(\bar{x}, \bar{y})$, whence $\exists \bar{y} \bigwedge_{i \in I_2} \gamma_i(\bar{x}, \bar{y}) \in \mathcal{A}(\Phi)(\bar{x})$. By assumption, every formula of $\mathcal{A}(\Phi)(\bar{x})$ is satisfiable modulo V ; then so are $\exists \bar{y} \bigwedge_{i \in I_2} \gamma_i(\bar{x}, \bar{y})$ and $\gamma(\bar{x}, \bar{y})$. \square

Proof of Proposition 5.2.16. It suffices to show just the ‘Only if’ direction of the result. Hence, consider a sentence Φ of $[\bigvee][\exists^* \bigwedge]^*$ FO given by $\Phi = \bigvee_{i \in I} \Psi_i$ where $\Psi_i \in [\exists^* \bigwedge]^*$ FO. Let $\mathcal{B} = \prod_{i \in I} \mathcal{A}(\Psi_i)$ where \prod denotes Cartesian product. We now show the following equivalences modulo V :

$$\Phi \leftrightarrow \bigvee_{i \in I} \bigwedge_{\gamma \in \mathcal{A}(\Psi_i)} \gamma \quad (5.1)$$

$$\leftrightarrow \bigwedge_{(\gamma_i) \in \mathcal{B}} \bigvee_{i \in I} \gamma_i \quad (5.2)$$

In equivalence Eq. 5.2 above, (γ_i) denotes a sequence in \mathcal{B} . Let $\mathcal{P}_{fin}(I)$ be the set of all finite subsets of I . We finally show the existence of a function $g : \mathcal{B} \rightarrow \mathcal{P}_{fin}(I)$ that gives the following equivalence

$$\Phi \leftrightarrow \bigwedge_{(\gamma_i) \in \mathcal{B}} \bigvee_{j \in g((\gamma_i))} \gamma_j \quad (5.3)$$

Observe that each disjunction in the RHS of Eq. 5.3 is a sentence of $\mathcal{A}(\Phi)$. Observe also that instead of ranging over all of \mathcal{B} in the RHS of Eq. 5.3 above, we can range over only a countable subset of \mathcal{B} , since the number of FO sentences over a finite vocabulary is countable. We now show the above equivalences to complete the proof. The equivalence Eq. 5.2 is obtained by applying the standard distributivity laws for conjunctions and disjunctions, to the sentence in the RHS of Eq. 5.1.

Proof of Eq. 5.1: Let $\Gamma = \bigvee_{i \in I} \bigwedge_{\gamma \in \mathcal{A}(\Psi_i)} \gamma$. Let \mathfrak{A} be a model of V such that $\mathfrak{A} \models \Phi$. Then $\mathfrak{A} \models \Psi_i$ for some $i \in I$. By Lemma 5.2.21, we have $\mathfrak{A} \models \mathcal{A}(\Psi_i)$, whence $\mathfrak{A} \models \Gamma$. Thus Φ implies Γ modulo V . Towards the converse, let \mathfrak{A} be a model of V such that $\mathfrak{A} \models \Gamma$. Then $\mathfrak{A} \models \mathcal{A}(\Psi_i)$ for some $i \in I$. Let $\Psi = \bigwedge (\text{Th}(\mathfrak{A}) \cup \{\Psi_i\})$, where $\text{Th}(\mathfrak{A})$ denotes the theory of \mathfrak{A} . It is easy to see that $\mathfrak{A} \models \mathcal{A}(\Psi)$ because any sentence γ in $\mathcal{A}(\Psi)$ is given by either $\gamma = \bigwedge Z$

The following figure gives the picture of the (partial) substructural characterizations, under Hypothesis 5.2.2, or equivalently, Hypothesis 5.2.13. (cf. Figure 5.2).




Figure 5.3: (Partial) conditional characterizations of $PSC(k)$ and $PSC(\lambda)$ theories

Chapter 6

Directions for future work

We propose as a part of future work, various directions that naturally arise from our results presented thus far.

1. We would like to investigate what syntactic subclasses of FO theories correspond exactly to $PSC(k)$ and $PSC(\aleph_0)$ theories. As Theorem 5.2.1 shows, these syntactic classes must semantically be subclasses of Σ_2^0 theories. For $PSC(k)$ theories, in addition to verifying whether Hypothesis 5.2.2 is true, we would further like to investigate what syntactic subclass of theories of $\exists^k \forall^*$ sentences characterizes $PSC(k)$ theories, assuming Hypothesis 5.2.2 holds. A technique to identify the latter syntactic subclass is to examine the syntactic properties of the FO theories given by Proposition 5.2.16, and exploit the fact that these theories are obtained from the finite approximations of the infinitary sentences of $[\bigvee] [\exists^k \bigwedge] \Pi_1^0$.
2. As “converses” to the investigations above, and as analogues of the semantic characterizations of Π_2^0 theories and theories of $\forall^k \exists^*$ sentences by $PCE(\lambda)$ and $PCE(k)$ respectively (cf. Theorem 5.1.1), we would like to semantically characterize Σ_2^0 theories and theories of $\exists^k \forall^*$ sentences, in terms of properties akin to (though not the same as) $PSC(\lambda)$ and $PSC(k)$. Likewise, as an analogue of Proposition 5.1.4(2), we would like to investigate if there are $PSC(\aleph_0)$ theories that are not equivalent to (i) any PSC theory (ii) any theory of $\exists^k \forall^*$ sentences, for any $k \geq 0$.
3. It is conceivable that many semantic properties of FO theories have natural and intuitive descriptions/characterizations in infinitary logics (Lemma 5.2.15 gives one such example). Then, results like Proposition 5.2.16 can be seen as “compilers” (in the sense of compilers used in computer science), in that they give a means of translating a “high level” description – via infinitary sentences that are known to be equivalent to FO theories – to an equivalent “low

level” description – via FO theories. The latter FO theories are obtained from appropriately defined finite approximations of the infinitary sentences. It would therefore be useful to investigate other infinitary logics and their fragments for which such “compiler-results” can be established. An interesting logic to investigate in this regard would be $\mathcal{L}_{\omega_1, \omega}$, which is well-known to enjoy excellent model-theoretic properties despite compactness theorem not holding of it [47].

4. Our results give characterizations of Σ_2^0 and Π_2^0 sentences in which the number of quantifiers in the leading block is given. As natural generalizations of these results, we can ask for characterizations of Σ_n^0 and Π_n^0 sentences for each $n \geq 2$, where the numbers of quantifiers in all the n blocks are given, and further extend these characterizations to theories. It may be noted that the results in the literature characterize Σ_n^0 and Π_n^0 theories as a whole and do not provide the finer characterizations suggested here.

We conclude this part of the thesis by presenting our ideas (in progress) on the last future work mentioned above, and suggesting concrete directions for pursuing the latter.

Directions for finer characterizations of Σ_n^0 and Π_n^0

For $n \geq 2$, let $\Sigma_n^0(k, l_1, *, l_2, *, \dots)$ be the class of all Σ_n^0 formulae in which the quantifier prefix is such that the leading block of quantifiers has k quantifiers, the $(2i)^{\text{th}}$ block has l_i quantifiers for $i \geq 1$, and the $(2i + 1)^{\text{th}}$ block has zero or more quantifiers for $i \geq 1$. Analogously, define the subclass $\Pi_n^0(k, l_1, *, l_2, *, \dots)$ of Π_n^0 . For $n \geq 1$, let $\Sigma_n^0(l_1, *, l_2, *, \dots)$ denote the class of all formulae of $\Pi_{n+1}^0(0, l_1, *, l_2, *, \dots)$; likewise, let $\Pi_n^0(l_1, *, l_2, *, \dots)$ denote the class of all formulae of $\Sigma_{n+1}^0(0, l_1, *, l_2, *, \dots)$.

Given a structure \mathfrak{A} and a k -tuple \bar{a} of \mathfrak{A} , the $\Pi_n^0(l_1, *, l_2, *, \dots)$ -type of \bar{a} in \mathfrak{A} is the set of all $\Pi_n^0(l_1, *, l_2, *, \dots)$ formulae having free variables among x_1, \dots, x_k , that are true of \bar{a} in \mathfrak{A} . We say a structure \mathfrak{B} *realizes* the $\Pi_n^0(l_1, *, l_2, *, \dots)$ -type of \bar{a} in \mathfrak{A} , if there exists a k -tuple \bar{b} of \mathfrak{B} such that the $\Pi_n^0(l_1, *, l_2, *, \dots)$ -type of \bar{b} in \mathfrak{B} contains (as a subset) the $\Pi_n^0(l_1, *, l_2, *, \dots)$ -type of \bar{a} in \mathfrak{A} . We now present generalizations of the notions of k -ary cover, $PSC(k)$ and $PCE(k)$.

Definition 6.1 ($((n; k, l_1, *, l_2, *, \dots)$ -ary cover). A collection R of structures is said to be a $((n; k, l_1, *, l_2, *, \dots)$ -ary cover of a structure \mathfrak{A} if for every k -tuple \bar{a} from \mathfrak{A} , there exists a structure in R that realizes the $\Pi_{n-1}^0(l_1, *, l_2, *, \dots)$ -type of \bar{a} in \mathfrak{A} .

Remark 6.2. Note that in the definition above, no structure in R need be a substructure of \mathfrak{A} . This is in contrast with the notion of k -ary cover as presented in Definition 3.2.1, where if R_1

is a k -ary cover of a structure \mathfrak{A}_1 , then each structure of R_1 is necessarily a substructure of \mathfrak{A}_1 .

Definition 6.3. Let T and V be given theories.

1. We say T is $PSC(n; k, l_1, *, l_2, *, \dots)$ modulo V if for every model \mathfrak{A} of $T \cup V$, there exists a k -tuple \bar{a} from \mathfrak{A} such that any model of V that realizes the $\Pi_{n-1}^0(l_1, *, l_2, *, \dots)$ -type of \bar{a} in \mathfrak{A} , is also a model of T .
2. We say T is $PCE(n; k, l_1, *, l_2, *, \dots)$ modulo V if for every model \mathfrak{A} of V and every $(n; k, l_1, *, l_2, *, \dots)$ -ary cover R of \mathfrak{A} , if each structure of R is a model of $T \cup V$, then \mathfrak{A} is a model of T .

We say a sentence ϕ is $PSC(n; k, l_1, *, l_2, *, \dots)$, resp. $PCE(n; k, l_1, *, l_2, *, \dots)$, modulo V , if the theory $\{\phi\}$ is $PSC(n; k, l_1, *, l_2, *, \dots)$, resp. $PCE(n; k, l_1, *, l_2, *, \dots)$, modulo V .

If $\phi(\bar{x})$ and $T(\bar{x})$ are respectively a formula and a theory, each of whose free variables are among \bar{x} , then for a theory V , the notions of ‘ $\phi(\bar{x})$ is $PSC(n; k, l_1, *, l_2, *, \dots)$ modulo V ’, ‘ $\phi(\bar{x})$ is $PCE(n; k, l_1, *, l_2, *, \dots)$ modulo V ’, ‘ $T(\bar{x})$ is $PSC(n; k, l_1, *, l_2, *, \dots)$ modulo V ’ and ‘ $T(\bar{x})$ is $PCE(n; k, l_1, *, l_2, *, \dots)$ modulo V ’ are defined similar to corresponding notions for $PSC(k)$ and $PCE(k)$.

We can now show the following results analogous to Lemma 3.2.5 and Theorem 4.1.1.

Lemma 6.4. Let V be a given theory. A formula $\phi(\bar{x})$ is $PSC(n; k, l_1, *, l_2, *, \dots)$ modulo V iff $\neg\phi(\bar{x})$ is $PCE(n; k, l_1, *, l_2, *, \dots)$ modulo V .

Theorem 6.5. Given a theory V , each of the following holds.

1. A formula $\phi(\bar{x})$ is $PSC(n; k, l_1, *, l_2, *, \dots)$ modulo V iff $\phi(\bar{x})$ is equivalent modulo V to a finite disjunction of $\Sigma_n^0(k, l_1, *, l_2, *, \dots)$ formulae, all of whose free variables are among \bar{x} .
2. A formula $\phi(\bar{x})$ is $PCE(n; k, l_1, *, l_2, *, \dots)$ modulo V iff $\phi(\bar{x})$ is equivalent modulo V to a finite conjunction of $\Pi_n^0(k, l_1, *, l_2, *, \dots)$ formulae, all of whose free variables are among \bar{x} .
3. A theory $T(\bar{x})$ is $PCE(n; k, l_1, *, l_2, *, \dots)$ modulo V iff $T(\bar{x})$ is equivalent modulo V to a theory of $\Pi_n^0(k, l_1, *, l_2, *, \dots)$ formulae, all of whose free variables are among \bar{x} .

The above theorem answers in part, the question raised in the last future work mentioned above. It also gives new semantic characterizations of Σ_n^0 and Π_n^0 sentences, via the properties $PSC_n = \bigcup_{k, l_1, l_2, \dots \in \mathbb{N}} PSC(n; k, l_1, *, l_2, *, \dots)$ and $PCE_n = \bigcup_{k, l_1, l_2, \dots \in \mathbb{N}} PCE(n; k, l_1, *, l_2, *, \dots)$ respectively.

A natural direction for future work that is suggested by Theorem 6.5 is the investigation of suitable variants of $PSC(n; k, l_1, *, l_2, *, \dots)$ and $PCE(n; k, l_1, *, l_2, *, \dots)$ that respectively characterize $\Sigma_n^0(k, l_1, *, l_2, *, \dots)$ and $\Pi_n^0(k, l_1, *, l_2, *, \dots)$ formulae exactly; likewise, an investigation of whether $PSC(n; k, l_1, *, l_2, *, \dots)$ or some suitable variant of it, characterizes theories of $\Sigma_n^0(k, l_1, *, l_2, *, \dots)$ formulae. If these characterizations are not obtained in general, then we would like to get them at least under plausible hypotheses (cf. Hypothesis 5.2.2). Of course, the question of characterizing Σ_n^0 and Π_n^0 sentences and theories in which the numbers of quantifiers in all blocks are given, still remains largely. However, observe that for the case of $n = 2$, the properties $PSC(n; k, l)$, resp. $PCE(n; k, l)$, do give semantic characterizations of finite disjunctions of Σ_2^0 sentences, resp. finite conjunctions of Π_2^0 sentences, in which the numbers of quantifiers in the first and second blocks are given numbers k and l respectively. This, in a sense, gives a finer characterization of Σ_2^0 and Π_2^0 than the one given by Theorem 4.1.1. But, we note that $PSC(n; k, l)$ and $PCE(n; k, l)$ are not combinatorial in nature. We would therefore like to investigate whether these notions, and more generally, the notions of $PSC(n; k, l_1, *, l_2, *, \dots)$ and $PCE(n; k, l_1, *, l_2, *, \dots)$ have combinatorial, or even purely algebraic equivalents, so that the “syntactical flavour” currently in the definitions of these notions, is eliminated.

Part II

Finite Model Theory

Chapter 7

Background and preliminaries

In this part of the thesis, we consider only finite structures over finite vocabularies τ that are relational, i.e. vocabularies that do not contain any constant or function symbols, unless explicitly stated otherwise. All the classes of structures that we consider are thus classes of finite relational structures. We denote classes of structures by \mathcal{S} possibly with numbers as subscripts, and assume these to be *closed under isomorphisms*. We consider two logics in this part of the thesis, one FO, and the other, an extension of it called *monadic second order logic*, denoted MSO. We use the notation \mathcal{L} to mean either FO or MSO. Any *notion or result stated for \mathcal{L}* means that the notion or result is stated for both FO and MSO. The classic references for the background from finite model theory presented in this chapter are [23, 34, 54].

7.1 Syntax and semantics of MSO

Syntax: The syntax of MSO extends that of FO by using MSO variables that range over subsets of the universes of structures, and using quantification (existential and universal) over these variables. We denote MSO variables using the capital letter X , possibly with numbers as subscripts. A sequence of MSO variables is denoted as \bar{X} . For a vocabulary τ , the notions of MSO terms and MSO formulae, and their free variables, are defined as follows.

1. An MSO term over τ is an FO term over τ , i.e. either a constant or an FO variable. A term that is a variable x has only one free variable, namely x . A constant has no free variables.
2. An atomic MSO formula over τ is either of the following.
 - An atomic FO formula over τ , i.e. either $t_1 = t_2$ or $R(t_1, \dots, t_n)$, where t_1, \dots, t_n are MSO terms over τ , and R is a relation symbol of τ of arity n . The free variables of

these are all FO variables, and are as defined in Section 2.1. These formulae have no free MSO variables.

- $X(t)$ where X is an MSO variable and t is an MSO term over τ . This formula has at most one free FO variable, namely the free variable of t (if any), and has exactly one free MSO variable, namely X .

3. Boolean combinations of MSO formulae using the boolean connectives \wedge , \vee and \neg are MSO formulae. The free variables of such formulae are defined analogously to those of boolean combinations of FO formulae (see Section 2.1).
4. Given an MSO formula φ , the formulae $\exists x\varphi$, $\forall x\varphi$, $\exists X\varphi$ and $\forall X\varphi$ are all MSO formulae. The free FO variables of $\exists x\varphi$ and $\forall x\varphi$ are the free FO variables of φ , except for x . The free MSO variables of these formulae are exactly those of φ . The free FO variables of $\exists X\varphi$ and $\forall X\varphi$ are exactly the free FO variables of φ , while the free MSO variables of these formulae are exactly those of φ , except for X .

We denote an MSO formula φ with free FO variables among \bar{x} , and free MSO variables among \bar{X} as $\varphi(\bar{x}, \bar{X})$. For φ as just mentioned, if \bar{X} is empty, i.e. if φ has no free MSO variables, then we denote φ as $\varphi(\bar{x})$. Like FO formulae, an MSO formula with no free variables is called a *sentence*, and an MSO formula with no quantifiers is called *quantifier-free*. Again, like for FO formulae, we denote MSO formulae using the Greek letters $\phi, \varphi, \psi, \chi, \xi, \gamma, \alpha$ or β .

Before looking at the semantics, we define the important notion of *quantifier-rank*, or simply *rank*, of an MSO formula φ , denoted $\text{rank}(\varphi)$. The definition is by structural induction.

1. If φ is quantifier-free, then $\text{rank}(\varphi) = 0$.
2. If $\varphi = \varphi_1 \wedge \varphi_2$, then $\text{rank}(\varphi) = \max(\text{rank}(\varphi_1), \text{rank}(\varphi_2))$. The same holds if $\varphi = \varphi_1 \vee \varphi_2$.
3. If $\varphi = \neg\varphi_1$, then $\text{rank}(\varphi) = \text{rank}(\varphi_1)$.
4. If φ is any of $\exists x\varphi_1$, $\forall x\varphi_1$, $\exists X\varphi_1$ or $\forall X\varphi_1$, then $\text{rank}(\varphi) = 1 + \text{rank}(\varphi_1)$.

The above definition also defines the rank of an FO formula, since every FO formula is also an MSO formula.

Remark 7.1.1. The MSO formulae φ that we consider in this part of the thesis always have *only FO free variables*, and *no MSO free variables* (although of course, the MSO formulae that build up φ surely would contain free MSO variables).

Since the semantics of MSO is defined inductively, we consider, only for the purposes of defining the semantics, MSO formulae with free MSO variables, in addition to free FO variables.

Semantics: Given a τ -structure \mathfrak{A} and an MSO formula $\varphi(\bar{x}, \bar{X})$, we define the notion of truth of $\varphi(\bar{x}, \bar{X})$ for a given assignment \bar{a} of elements of \mathfrak{A} , to \bar{x} and a given assignment \bar{A} of subsets of elements of \mathfrak{A} , to \bar{X} . We denote by $(\mathfrak{A}, \bar{a}, \bar{A}) \models \varphi(\bar{x}, \bar{X})$, that $\varphi(\bar{x}, \bar{X})$ is true in \mathfrak{A} for the assignments \bar{a} to \bar{x} and \bar{A} to \bar{X} , and call $(\mathfrak{A}, \bar{a}, \bar{A})$ a model of $\varphi(\bar{x}, \bar{X})$. We give the semantics only for the syntactic features of MSO that are different from those of FO. Below, $\bar{X} = (X_1, \dots, X_n)$ and $\bar{A} = (A_1, \dots, A_n)$.

- If $\varphi(\bar{x}, \bar{X})$ is the formula $X_i(t)$, then $(\mathfrak{A}, \bar{a}, \bar{A}) \models \varphi(\bar{x}, \bar{X})$ iff $t^{\mathfrak{A}}(\bar{a}) \in A_i$.
- If $\varphi(\bar{x}, \bar{X})$ is the formula $\exists X_{n+1} \varphi_1(\bar{x}, \bar{X}, X_{n+1})$, then $(\mathfrak{A}, \bar{a}, \bar{A}) \models \varphi(\bar{x}, \bar{X})$ iff there exists $A_{n+1} \subseteq U_{\mathfrak{A}}$ such that $(\mathfrak{A}, \bar{a}, \bar{A}, A_{n+1}) \models \varphi_1(\bar{x}, \bar{X}, X_{n+1})$.
- If $\varphi(\bar{x}, \bar{X})$ is the formula $\forall X_{n+1} \varphi_1(\bar{x}, \bar{X}, X_{n+1})$, then $(\mathfrak{A}, \bar{a}, \bar{A}) \models \varphi(\bar{x}, \bar{X})$ iff for all $A_{n+1} \subseteq U_{\mathfrak{A}}$, it is the case that $(\mathfrak{A}, \bar{a}, \bar{A}, A_{n+1}) \models \varphi_1(\bar{x}, \bar{X}, X_{n+1})$.

If $\varphi(\bar{a})$ is an MSO formula with no MSO variables, then we denote the truth of $\varphi(\bar{x})$ in \mathfrak{A} for an assignment \bar{a} to \bar{x} as $(\mathfrak{A}, \bar{a}) \models \varphi(\bar{x})$. For a sentence φ , we denote the truth of φ in \mathfrak{A} as $\mathfrak{A} \models \varphi$.

7.2 Adaptations of classical model theory concepts to the finite model theory setting

We assume familiarity with the notions introduced in Chapter 2. We adapt some of these notions to versions of these *over* classes of structures. For other notions, they are exactly as defined in Chapter 2.

1. Consistency, validity, entailment, and equivalence *over* classes of structures

Given a non-empty class \mathcal{S} of structures, we say a formula $\varphi(\bar{x})$ is *satisfiable over* \mathcal{S} if there exists structure $\mathfrak{A} \in \mathcal{S}$ and a tuple \bar{a} of \mathfrak{A} such that $|\bar{a}| = |\bar{x}|$ and $(\mathfrak{A}, \bar{a}) \models \varphi(\bar{x})$. We say $\varphi(\bar{x})$ is *unsatisfiable over* \mathcal{S} if $\varphi(\bar{x})$ is not satisfiable over \mathcal{S} . We say $\psi(\bar{x})$ *entails* $\varphi(\bar{x})$ *over* \mathcal{S} if any model (\mathfrak{A}, \bar{a}) of $\psi(\bar{x})$ such that $\mathfrak{A} \in \mathcal{S}$ is also a model of $\varphi(\bar{x})$. We say $\psi(\bar{x})$ and $\varphi(\bar{x})$ are *equivalent over* \mathcal{S} , or simply *\mathcal{S} -equivalent*, if $\psi(\bar{x})$ entails $\varphi(\bar{x})$ over \mathcal{S} , and vice-versa.

2. (m, \mathcal{L}) -types: Given $m \in \mathbb{N}$, a τ -structure \mathfrak{A} and a k -tuple \bar{a} from \mathfrak{A} , the (m, \mathcal{L}) -type of \bar{a} in \mathfrak{A} , denoted $\text{tp}_{\mathfrak{A}, \bar{a}, m, \mathcal{L}}(x_1, \dots, x_k)$, is the set of all $\mathcal{L}(\tau)$ formulae of rank at most m , whose free variables are among x_1, \dots, x_k , and that are true of \bar{a} in \mathfrak{A} . We denote by $\text{Th}_{m, \mathcal{L}}(\mathfrak{A})$, the set of all $\mathcal{L}(\tau)$ sentences of rank at most m that are true in \mathfrak{A} . Given a τ -structure \mathfrak{B} , we say that \mathfrak{A} and \mathfrak{B} are (m, \mathcal{L}) -equivalent, denoted $\mathfrak{A} \equiv_{m, \mathcal{L}} \mathfrak{B}$ if $\text{Th}_{m, \mathcal{L}}(\mathfrak{A}) = \text{Th}_{m, \mathcal{L}}(\mathfrak{B})$. Observe that if \bar{a} and

\bar{a} are k -tuples from \mathfrak{A} and \mathfrak{B} respectively, then $\text{tp}_{\mathfrak{A}, \bar{a}, m, \mathcal{L}}(x_1, \dots, x_k) = \text{tp}_{\mathfrak{B}, \bar{b}, m, \mathcal{L}}(x_1, \dots, x_k)$ iff $(\mathfrak{A}, \bar{a}) \equiv_{m, \mathcal{L}} (\mathfrak{B}, \bar{b})$. If $\mathcal{L} = \text{FO}$, then we also denote $\equiv_{m, \text{FO}}$ simply as \equiv_m , following standard notation in the literature. It is easy to see that $\equiv_{m, \mathcal{L}}$ is an equivalence relation over all structures. Following is an important result concerning the $\equiv_{m, \mathcal{L}}$ relation.

Proposition 7.2.1 (Proposition 7.5, ref. [54]). *There exists a computable function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that for each $m \in \mathbb{N}$, the index of the $\equiv_{m, \mathcal{L}}$ relation is at most $f(m)$.*

Given a class \mathcal{S} of structures and $m \in \mathbb{N}$, we let $\Delta_{\mathcal{L}}(m, \mathcal{S})$ denote the set of all equivalence classes of the $\equiv_{m, \mathcal{L}}$ relation restricted to the structures in \mathcal{S} . We denote by $\Lambda_{\mathcal{S}, \mathcal{L}} : \mathbb{N} \rightarrow \mathbb{N}$ a fixed computable function with the property that $\Lambda_{\mathcal{S}, \mathcal{L}}(m) \geq |\Delta_{\mathcal{L}}(m, \mathcal{S})|$. The existence of $\Lambda_{\mathcal{S}, \mathcal{L}}$ is guaranteed by Proposition 7.2.1.

The notion of $\equiv_{m, \mathcal{L}}$ has a characterization in terms of *Ehrenfeucht-Fr iss e* games for \mathcal{L} . We describe these in the next section.

7.3 \mathcal{L} -Ehrenfeucht-Fr iss e games

We first define the notion of *partial isomorphism* that is crucially used in the definition of the games we present below. We assume for this section, that the vocabulary τ possibly contains constant symbols, in addition to relation symbols. Let $\mathfrak{A}, \mathfrak{B}$ be two τ -structures, and $\bar{a} = (a_1, \dots, a_m)$ and $\bar{b} = (b_1, \dots, b_m)$ be two m -tuples from \mathfrak{A} and \mathfrak{B} respectively. Then (\bar{a}, \bar{b}) defines a partial isomorphism between \mathfrak{A} and \mathfrak{B} if the following are true.

1. For $1 \leq i, j \leq m$, we have $a_i = a_j$ iff $b_i = b_j$.
2. For every constant symbol $c \in \tau$ and $1 \leq i \leq m$, we have $a_i = c^{\mathfrak{A}}$ iff $b_i = c^{\mathfrak{B}}$.
3. For every r -ary relation symbol $R \in \tau$ and every sequence (i_1, \dots, i_r) of numbers from $\{1, \dots, m\}$, we have $(a_{i_1}, \dots, a_{i_r}) \in R^{\mathfrak{A}}$ iff $(b_{i_1}, \dots, b_{i_r}) \in R^{\mathfrak{B}}$.

If τ contains no constant symbols, then the map $a_i \mapsto b_i$ for $1 \leq i \leq m$ is an isomorphism from the substructure of \mathfrak{A} induced by $\{a_1, \dots, a_m\}$ to the substructure of \mathfrak{B} induced by $\{b_1, \dots, b_m\}$.

The FO-Ehrenfeucht-Fr iss e game

The FO-Ehrenfeucht-Fr iss e game, or simply the FO-EF game, is played on two given structures \mathfrak{A} and \mathfrak{B} , and by two players, called *spoiler* and *duplicator*. The spoiler tries to show that the two structures are non-isomorphic, while the duplicator tries to show otherwise. The FO-EF game of m rounds between \mathfrak{A} and \mathfrak{B} is as defined below. Each round consists of the following steps.

1. The spoiler chooses one of the structures and picks an element from it.
2. The duplicator responds by picking an element of the other structure.

At the end of m rounds, let $\bar{a} = (a_1, \dots, a_m)$ be the elements chosen from \mathfrak{A} and let $\bar{b} = (b_1, \dots, b_m)$ be the elements chosen from \mathfrak{B} . Let c_1, \dots, c_p be the constant symbols of τ , and let $\bar{c}^{\mathfrak{A}} = (c_1^{\mathfrak{A}}, \dots, c_p^{\mathfrak{A}})$ and $\bar{c}^{\mathfrak{B}} = (c_1^{\mathfrak{B}}, \dots, c_p^{\mathfrak{B}})$. Call (\bar{a}, \bar{b}) as a *play* of m rounds of the FO-EF game on \mathfrak{A} and \mathfrak{B} . The duplicator is said to *win* the play (\bar{a}, \bar{b}) iff $((\bar{a}, \bar{c}^{\mathfrak{A}}), (\bar{b}, \bar{c}^{\mathfrak{B}}))$ is a partial isomorphism between \mathfrak{A} and \mathfrak{B} . An *m -round strategy for the duplicator in the FO-EF game* on \mathfrak{A} and \mathfrak{B} is a function $\mathbf{S} : \bigcup_{i=0}^{m-1} ((U_{\mathfrak{A}} \times U_{\mathfrak{B}})^i \times (U_{\mathfrak{A}} \cup U_{\mathfrak{B}})) \rightarrow (U_{\mathfrak{A}} \cup U_{\mathfrak{B}})$ such that for all $i \in \{0, \dots, m-1\}$, for all $(a_1, b_1), \dots, (a_i, b_i) \in (U_{\mathfrak{A}} \times U_{\mathfrak{B}})$ and all $d \in (U_{\mathfrak{A}} \cup U_{\mathfrak{B}})$, it is the case that $\mathbf{S}((a_1, b_1), \dots, (a_i, b_i), d) \in U_{\mathfrak{A}}$ iff $d \in U_{\mathfrak{B}}$. The duplicator is said to have a *winning strategy* in the m -round FO-EF game on \mathfrak{A} and \mathfrak{B} if there exists an m -round strategy \mathbf{S} for the duplicator in the FO-EF game on \mathfrak{A} and \mathfrak{B} , such that the duplicator wins every play of m rounds of the FO-EF game on \mathfrak{A} and \mathfrak{B} , in which the duplicator responds in accordance with \mathbf{S} .

The following theorem shows that the existence of a winning strategy for the duplicator in the FO-EF game of m -rounds on \mathfrak{A} and \mathfrak{B} characterizes (m, FO) -equivalence of \mathfrak{A} and \mathfrak{B} (see Theorem 3.9 of [54]).

Theorem 7.3.1 (Ehrenfeucht-Fr  iss  ). *Let \mathfrak{A} and \mathfrak{B} be two structures over a vocabulary that possibly contains constant symbols. Let \bar{a} and \bar{b} be given tuples of elements from \mathfrak{A} and \mathfrak{B} respectively. Then $(\mathfrak{A}, \bar{a}) \equiv_{m, \text{FO}} (\mathfrak{B}, \bar{b})$ iff the duplicator has a winning strategy in the m -round FO-EF game on (\mathfrak{A}, \bar{a}) and (\mathfrak{B}, \bar{b}) .*

The MSO-Ehrenfeucht-Fr  iss   game

The MSO-Ehrenfeucht-Fr  iss   game, or simply the MSO-EF game, is similar to the FO-EF game. It is played on two given structures \mathfrak{A} and \mathfrak{B} , and by two players, namely the spoiler and duplicator. The difference with the FO-EF game is that in an MSO-EF game of m rounds on \mathfrak{A} and \mathfrak{B} , in each round there are two kinds of moves.

1. Point move: This is like in the FO-EF game on \mathfrak{A} and \mathfrak{B} . The spoiler chooses one of the structures and picks an element from it. The duplicator responds by picking an element of the other structure.
2. Set move: The spoiler chooses a subset of elements from one of the structures. The duplicator responds by picking a subset of elements of the other structure.

At the end of m rounds, let $\bar{a} = (a_1, \dots, a_p)$ be the elements chosen from \mathfrak{A} and let $\bar{b} = (b_1, \dots, b_p)$ be the elements chosen from \mathfrak{B} . Likewise, let $\bar{A} = (A_1, \dots, A_r)$ be the sets cho-

sen from \mathfrak{A} and $\bar{B} = (B_1, \dots, B_r)$ be the sets chosen from \mathfrak{B} such that $p + r = m$. Let c_1, \dots, c_s be the constant symbols of τ , and let $\bar{c}^{\mathfrak{A}} = (c_1^{\mathfrak{A}}, \dots, c_s^{\mathfrak{A}})$ and $\bar{c}^{\mathfrak{B}} = (c_1^{\mathfrak{B}}, \dots, c_s^{\mathfrak{B}})$. Call $((\bar{a}, \bar{A}), (\bar{b}, \bar{B}))$ as a play of m rounds of the MSO-EF game on \mathfrak{A} and \mathfrak{B} . The duplicator wins the play $((\bar{a}, \bar{A}), (\bar{b}, \bar{B}))$ if $((\bar{a}, \bar{c}^{\mathfrak{A}}), (\bar{b}, \bar{c}^{\mathfrak{B}}))$ is a partial isomorphism between (\mathfrak{A}, \bar{A}) and (\mathfrak{B}, \bar{B}) . Note that the latter implies $a_i \in A_j$ iff $b_i \in B_j$, for $i \in \{1, \dots, p\}$ and $j \in \{1, \dots, r\}$. We now define the notion of a strategy for the duplicator in an MSO-EF game analogous to that for an FO-EF game. For a set X , let 2^X denote the powerset of X . An m -round strategy for the duplicator in the MSO-EF game on \mathfrak{A} and \mathfrak{B} is a function $\mathbf{S} : \bigcup_{i=0}^{m-1} ((\mathcal{U}_{\mathfrak{A}} \cup 2^{\mathcal{U}_{\mathfrak{A}}}) \times (\mathcal{U}_{\mathfrak{B}} \cup 2^{\mathcal{U}_{\mathfrak{B}}}))^i \times (\mathcal{U}_{\mathfrak{A}} \cup 2^{\mathcal{U}_{\mathfrak{A}}} \cup \mathcal{U}_{\mathfrak{B}} \cup 2^{\mathcal{U}_{\mathfrak{B}}}) \rightarrow (\mathcal{U}_{\mathfrak{A}} \cup 2^{\mathcal{U}_{\mathfrak{A}}} \cup \mathcal{U}_{\mathfrak{B}} \cup 2^{\mathcal{U}_{\mathfrak{B}}})$ such that for all $i \in \{0, \dots, m-1\}$, for all $(d_1, e_1), \dots, (d_i, e_i) \in ((\mathcal{U}_{\mathfrak{A}} \cup 2^{\mathcal{U}_{\mathfrak{A}}}) \times (\mathcal{U}_{\mathfrak{B}} \cup 2^{\mathcal{U}_{\mathfrak{B}}}))$ and all $d \in (\mathcal{U}_{\mathfrak{A}} \cup 2^{\mathcal{U}_{\mathfrak{A}}} \cup \mathcal{U}_{\mathfrak{B}} \cup 2^{\mathcal{U}_{\mathfrak{B}}})$, it is the case that (i) $\mathbf{S}((d_1, e_1), \dots, (d_i, e_i), d) \in (\mathcal{U}_{\mathfrak{A}} \cup \mathcal{U}_{\mathfrak{B}})$ iff $d \in (\mathcal{U}_{\mathfrak{A}} \cup \mathcal{U}_{\mathfrak{B}})$, (ii) $\mathbf{S}((d_1, e_1), \dots, (d_i, e_i), d) \in \mathcal{U}_{\mathfrak{A}}$ iff $d \in \mathcal{U}_{\mathfrak{B}}$, and (iii) $\mathbf{S}((d_1, e_1), \dots, (d_i, e_i), d) \in 2^{\mathcal{U}_{\mathfrak{A}}}$ iff $d \in 2^{\mathcal{U}_{\mathfrak{B}}}$. The duplicator is said to have a *winning strategy* in the m -round MSO-EF game on \mathfrak{A} and \mathfrak{B} if there exists an m -round strategy \mathbf{S} for the duplicator in the MSO-EF game on \mathfrak{A} and \mathfrak{B} such that the duplicator wins every play of m rounds of the MSO-EF game on \mathfrak{A} and \mathfrak{B} , when the duplicator responds in accordance with \mathbf{S} .

Like Theorem 7.3.1, the following theorem shows that the existence of a winning strategy for the duplicator in the MSO-EF game of m -rounds on \mathfrak{A} and \mathfrak{B} characterizes (m, MSO) -equivalence of \mathfrak{A} and \mathfrak{B} . This theorem is stated more generally in Theorem 7.7 of [54]. However, we need the theorem only as stated below.

Theorem 7.3.2. *Let \mathfrak{A} and \mathfrak{B} be two structures over a vocabulary that possibly contains constant symbols. Let \bar{a} and \bar{b} be given tuples of elements from \mathfrak{A} and \mathfrak{B} respectively. Then $(\mathfrak{A}, \bar{a}) \equiv_{m, \text{MSO}} (\mathfrak{B}, \bar{b})$ iff the duplicator has a winning strategy in the m -round MSO-EF game on (\mathfrak{A}, \bar{a}) and (\mathfrak{B}, \bar{b}) .*

7.4 Translation schemes

We recall the notion of translation schemes from the literature [57]. These were first introduced in the context of classical model theory, and are known in the literature by different names, like *FO interpretations*, *transductions*, etc. We define these below and look at some of their properties subsequently.

Let τ and σ be given vocabularies, and $t \geq 1$ be a natural number. Let \bar{x}_0 be a fixed t -tuple of

first order variables, and for each relation $R \in \sigma$ of arity $\#R$, let \bar{x}_R be a fixed $(t \times \#R)$ -tuple of first order variables. A $(t, \tau, \sigma, \mathcal{L})$ -translation scheme $\Xi = (\xi, (\xi_R)_{R \in \sigma})$ is a sequence of formulas of $\mathcal{L}(\tau)$ such that the free variables of ξ are among those in \bar{x}_0 , and for $R \in \sigma$, the free variables of ξ_R are among those in \bar{x}_R . When t, σ and τ are clear from context, we call Ξ simply as a translation scheme. We call t as the *dimension* of Ξ . If $t = 1$, we say Ξ is a *scalar* translation scheme, and if $t \geq 2$, we say Ξ is a *vectorized* translation scheme. In our results in the subsequent chapters, we consider vectorized translation schemes when $\mathcal{L} = \text{FO}$, and scalar translation schemes when $\mathcal{L} = \text{MSO}$.

One can associate with a translation scheme Ξ , two partial maps: (i) Ξ^* from τ -structures to σ -structures (ii) Ξ^\sharp from $\mathcal{L}(\sigma)$ formulae to $\mathcal{L}(\tau)$ formulae, each of which we define below. For the ease of readability, we abuse notation slightly and use Ξ to denote both Ξ^* and Ξ^\sharp .

1. Given a τ -structure \mathfrak{A} , the σ -structure $\Xi(\mathfrak{A})$ is defined as follows.

1. $U_{\Xi(\mathfrak{A})} = \{\bar{a} \in U_{\mathfrak{A}}^t \mid (\mathfrak{A}, \bar{a}) \models \xi(\bar{x}_0)\}$.
2. For each $R \in \sigma$ of arity n , the interpretation of R in $\Xi(\mathfrak{A})$ is the set $\{\bar{a} \in U_{\mathfrak{A}}^{t \times n} \mid (\mathfrak{A}, \bar{a}) \models \xi_R(\bar{x}_R)\}$.

2. We define the map Ξ from $\mathcal{L}(\sigma)$ formulae to $\mathcal{L}(\tau)$ formulae. We first define this map for the case of a $(t, \tau, \sigma, \text{FO})$ -translation scheme. Given a $\text{FO}(\sigma)$ formula $\varphi(\bar{x})$ where $\bar{x} = (x_1, \dots, x_n)$, the $\text{FO}(\tau)$ formula $\Xi(\varphi)(\bar{x}_1, \dots, \bar{x}_n)$ where $\bar{x}_i = (x_{i,1}, \dots, x_{i,t})$ for $1 \leq i \leq n$, is as defined below.

1. If $\varphi(\bar{x})$ is the formula $R(x_1, \dots, x_r)$ for an r -ary relation symbol $R \in \sigma$, then

$$\Xi(\varphi)(\bar{x}_1, \dots, \bar{x}_n) = \xi_R(\bar{x}_1, \dots, \bar{x}_n) \wedge \bigwedge_{i=1}^{i=r} \xi(x_{i,1}, \dots, x_{i,t})$$

2. If $\varphi(\bar{x})$ is the formula $x_1 = x_2$, then

$$\Xi(\varphi)(\bar{x}_1, \dots, \bar{x}_n) = \bigwedge_{j=1}^{j=t} (x_{1,j} = x_{2,j}) \wedge \bigwedge_{i=1}^{i=2} \xi(x_{i,1}, \dots, x_{i,t})$$

3. If $\varphi(\bar{x}) = \varphi_1(\bar{x}) \wedge \varphi_2(\bar{x})$, then

$$\Xi(\varphi)(\bar{x}_1, \dots, \bar{x}_n) = \Xi(\varphi_1)(\bar{x}_1, \dots, \bar{x}_n) \wedge \Xi(\varphi_2)(\bar{x}_1, \dots, \bar{x}_n)$$

The same holds with \wedge replaced with \vee .

4. If $\varphi(\bar{x}) = \neg \varphi_1(\bar{x})$, then

$$\Xi(\varphi)(\bar{x}_1, \dots, \bar{x}_n) = \neg \Xi(\varphi_1)(\bar{x}_1, \dots, \bar{x}_n)$$

5. If $\varphi(\bar{x}) = \exists y \varphi_1(\bar{x}, y)$, then for $\bar{y} = (y_1, \dots, y_t)$

$$\Xi(\varphi)(\bar{x}_1, \dots, \bar{x}_n) = \exists \bar{y} (\Xi(\varphi_1)(\bar{x}_1, \dots, \bar{x}_n, \bar{y}) \wedge \xi(y_1, \dots, y_t))$$

6. If $\varphi(\bar{x}) = \forall y \varphi_1(\bar{x}, y)$, then for $\bar{y} = (y_1, \dots, y_t)$

$$\Xi(\varphi)(\bar{x}_1, \dots, \bar{x}_n) = \forall \bar{y} (\xi(y_1, \dots, y_t) \rightarrow \Xi(\varphi_1)(\bar{x}_1, \dots, \bar{x}_n, \bar{y}))$$

We now define Ξ for a $(1, \tau, \sigma, \text{MSO})$ -translation scheme. Given an $\text{MSO}(\sigma)$ formula $\varphi(\bar{x})$ where $\bar{x} = (x_1, \dots, x_n)$, the $\text{MSO}(\tau)$ formula $\Xi(\varphi)(\bar{x})$, is as defined below.

1. If $\varphi(\bar{x})$ is an FO atomic formula, then $\Xi(\varphi)(\bar{x})$ is as defined in the case of FO above.
2. If $\varphi(\bar{x})$ is the formula $X(x_1)$ for an MSO variable X ,

$$\Xi(\varphi)(\bar{x}) = X(x_1) \wedge \xi(x_1)$$

3. For boolean combinations, and quantification over FO variables, $\Xi(\varphi)(\bar{x})$ is as defined in the case of FO above.
4. If $\varphi(\bar{x}) = \exists Y \varphi_1(\bar{x}, Y)$, then

$$\Xi(\varphi)(\bar{x}) = \exists Y \Xi(\varphi_1)(\bar{x}, Y)$$

5. If $\varphi(\bar{x}) = \forall Y \varphi_1(\bar{x}, Y)$, then

$$\Xi(\varphi)(\bar{x}) = \forall Y \Xi(\varphi_1)(\bar{x}, Y)$$

Properties of translation schemes:

For a $(t, \tau, \sigma, \mathcal{L})$ -translation scheme $\Xi = (\xi, (\xi_R)_{R \in \sigma})$, let $\text{rank}(\Xi)$ denote the maximum of the quantifier ranks of the formulae ξ and ξ_R for each $R \in \sigma$.

Lemma 7.4.1. *Let Ξ be a $(t, \tau, \sigma, \mathcal{L})$ -translation scheme, and let φ be an $\mathcal{L}(\sigma)$ formula of quantifier rank m .*

1. *If $\mathcal{L} = \text{FO}$, then $\Xi(\varphi)$ is an $\text{FO}(\tau)$ formula having quantifier rank at most $t \cdot m + \text{rank}(\Xi)$.*
2. *If $\mathcal{L} = \text{MSO}$ and Ξ is scalar, then $\Xi(\varphi)$ is an $\text{MSO}(\tau)$ formula having quantifier rank at most $m + \text{rank}(\Xi)$.*

The following proposition relates the application of transductions to structures and formulas to each other.

Proposition 7.4.2. *Let Ξ be either a (t, τ, σ, FO) -translation scheme, or a $(1, \tau, \sigma, MSO)$ -translation scheme. Then for every $\mathcal{L}(\sigma)$ formula $\varphi(x_1, \dots, x_n)$ where $n \geq 0$, for every τ -structure \mathfrak{A} and for every n -tuple $(\bar{a}_1, \dots, \bar{a}_n)$ from $\Xi(\mathfrak{A})$, the following holds.*

$$\begin{aligned} (\Xi(\mathfrak{A}), \bar{a}_1, \dots, \bar{a}_n) &\models \varphi(x_1, \dots, x_n) \\ \text{iff} \quad (\mathfrak{A}, \bar{a}_1, \dots, \bar{a}_n) &\models \Xi(\varphi)(\bar{x}_1, \dots, \bar{x}_n) \end{aligned}$$

where $\bar{x}_i = (x_{i,1}, \dots, x_{i,t})$ for each $i \in \{1, \dots, n\}$.

An immediate consequence of Lemma 7.4.1 and Proposition 7.4.2 is the following.

Corollary 7.4.3. *Let Ξ be a $(t, \tau, \sigma, \mathcal{L})$ -translation scheme. Let $m, r \in \mathbb{N}$ be such that $r = t \cdot m + \text{rank}(\Xi)$. Suppose \mathfrak{A} and \mathfrak{B} are τ -structures, and suppose $\bar{a}_1, \dots, \bar{a}_n$, resp. $\bar{b}_1, \dots, \bar{b}_n$, are n elements from $\Xi(\mathfrak{A})$, resp. $\Xi(\mathfrak{B})$.*

1. *If $(\mathfrak{A}, \bar{a}_1, \dots, \bar{a}_n) \equiv_{r, FO} (\mathfrak{B}, \bar{b}_1, \dots, \bar{b}_n)$, then $(\Xi(\mathfrak{A}), \bar{a}_1, \dots, \bar{a}_n) \equiv_{m, FO} (\Xi(\mathfrak{B}), \bar{b}_1, \dots, \bar{b}_n)$.*
2. *If $(\mathfrak{A}, \bar{a}_1, \dots, \bar{a}_n) \equiv_{r, MSO} (\mathfrak{B}, \bar{b}_1, \dots, \bar{b}_n)$, then $(\Xi(\mathfrak{A}), \bar{a}_1, \dots, \bar{a}_n) \equiv_{m, MSO} (\Xi(\mathfrak{B}), \bar{b}_1, \dots, \bar{b}_n)$, when Ξ is scalar.*

We use the above properties of translation schemes in our results in the forthcoming chapters.

Before we move to the next chapter, we define three notions that will frequently appear in our discussions. These are the notions of hereditariness, disjoint union and cartesian product. Given a class \mathcal{S}_1 of structures and a subclass \mathcal{S}_2 of \mathcal{S}_1 , we say \mathcal{S}_2 is *hereditary over \mathcal{S}_1* , if \mathcal{S}_2 is *PS* over \mathcal{S}_1 (see Section 2.5 for the definition of *PS*). A class \mathcal{S} is *hereditary* if it is hereditary over the class of all (finite) structures. Given two τ -structures \mathfrak{A} and \mathfrak{B} , the *disjoint union* of \mathfrak{A} and \mathfrak{B} , denoted $\mathfrak{A} \sqcup \mathfrak{B}$, is defined as follows. Let \mathfrak{B}' be an isomorphic copy of \mathfrak{B} whose universe is disjoint with that of \mathfrak{A} . Then $\mathfrak{A} \sqcup \mathfrak{B}$ is defined upto isomorphism as the structure \mathcal{C} such that (i) $U_{\mathcal{C}} = U_{\mathfrak{A}} \cup U_{\mathfrak{B}'}$ and (ii) $R^{\mathcal{C}} = R^{\mathfrak{A}} \cup R^{\mathfrak{B}'}$ for each relation symbol $R \in \tau$. Finally, the *cartesian product* of \mathfrak{A} and \mathfrak{B} , denoted $\mathfrak{A} \times \mathfrak{B}$, is the structure \mathcal{C} such that (i) $U_{\mathcal{C}} = U_{\mathfrak{A}} \times U_{\mathfrak{B}}$ and (ii) for each n -ary relation symbol $R \in \tau$, for each n -tuple $((a_1, b_1), \dots, (a_n, b_n))$ from \mathcal{C} , where (a_1, \dots, a_n) is an n -tuple from \mathfrak{A} and (b_1, \dots, b_n) is an n -tuple from \mathfrak{B} , we have $((a_1, b_1), \dots, (a_n, b_n)) \in R^{\mathcal{C}}$ iff $(a_1, \dots, a_n) \in R^{\mathfrak{A}}$ and $(b_1, \dots, b_n) \in R^{\mathfrak{B}}$.

Chapter 8

The need to investigate new classes of finite structures for $\text{GLT}(k)$

As already mentioned in the introduction, the class of finite structures behaves very differently compared to the class of arbitrary structures. The failure of the compactness theorem in the finite causes a collapse of the proofs of most of the classical preservation theorems in the finite. Worse still, the statements of these theorems also fail in the finite. The Łoś-Tarski theorem is an instance of these failures [37, 68, 81].

Theorem 8.1 (Tait 1959, Gurevich-Shelah 1984). *There is a sentence that is preserved under substructures over the class of all finite structures, but that is not equivalent, over all finite structures, to any Π_1^0 sentence.*

In the last 15 years, a lot of research in the area of preservation theorems in finite model theory has focussed on identifying classes of finite structures over which classical preservation theorems hold. The following theorem from [7] identifies classes of finite structures that satisfy structural restrictions that are interesting from a computational standpoint and also from the standpoint of modern graph structure theory, and shows that these classes are “well-behaved” with respect some classical preservation theorems, indeed in particular the Łoś-Tarski theorem.

Theorem 8.2 (Atserias-Dawar-Grohe, 2008). *The Łoś-Tarski theorem holds over each of the following classes of finite structures:*

1. *Any class of acyclic structures that is closed under substructures and disjoint unions.*

2. Any class of bounded degree structures that is closed under substructures and disjoint unions.
3. The class of all structures of tree-width at most k , for each $k \in \mathbb{N}$.

It is natural to ask what happens to $\text{GLT}(k)$ over the classes of finite structures referred to above. Unfortunately, $\text{GLT}(k)$ fails in general over each of these classes. We show this failure first for the class of all finite structures, and then for the special classes of finite structures considered in Theorem 8.2.

8.1 Failure of $\text{GLT}(k)$ over all finite structures

The failure of the Łoś-Tarski theorem over the class of all finite structures already implies the failure of $\text{GLT}(0)$ over this class. We show below that this failure happens for $\text{GLT}(k)$ for each $k \geq 0$. In fact, we show something stronger.

Proposition 8.1.1 (Failure of $\text{GLT}(k)$ in the finite). *There exists a vocabulary τ such that if \mathcal{S} is the class of all finite τ -structures, then for each $k \geq 0$, there exists an $\text{FO}(\tau)$ sentence ψ_k that is preserved under substructures over \mathcal{S} , but that is not \mathcal{S} -equivalent to any $\exists^k \forall^*$ sentence. It follows that there is a sentence that is $\text{PSC}(k)$ over \mathcal{S} (ψ_k being one such sentence) but that is not \mathcal{S} -equivalent to any $\exists^k \forall^*$ sentence.*

Proof. The second part of the proposition follows from the first part since a sentence that is preserved under substructures over \mathcal{S} is also $\text{PSC}(k)$ over \mathcal{S} for each $k \geq 0$. We now prove the first part of the proposition.

Consider the vocabulary $\tau = \{\leq, S, P, c, d\}$ where \leq and S are both binary relation symbols, P is a unary relation symbol, and c and d are constant symbols. The sentence ψ_k is constructed along the lines of the known counterexample to the Łoś-Tarski theorem in the finite. Following are the details.

$$\begin{aligned}
 \psi_k &= (\xi_1 \wedge \xi_2 \wedge \xi_3) \wedge \neg(\xi_4 \wedge \xi_5) \\
 \xi_1 &= \text{“} \leq \text{ is a linear order ”} \\
 \xi_2 &= \text{“} c \text{ is minimum under } \leq \text{ and } d \text{ is maximum under } \leq \text{”} \\
 \xi_3 &= \forall x \forall y S(x, y) \rightarrow \text{“} y \text{ is the successor of } x \text{ under } \leq \text{”} \\
 \xi_4 &= \forall x (x \neq d) \rightarrow \exists y S(x, y) \\
 \xi_5 &= \text{“} \text{There exist at most } k \text{ elements in (the set interpreting) } P \text{”}
 \end{aligned}$$

It is easy to see that each of ξ_1, ξ_2, ξ_3 and ξ_5 can be expressed using a universal sentence. In particular, ξ_1 and ξ_3 can be expressed using a \forall^3 sentence each, ξ_2 can be expressed using a \forall sentence, and ξ_5 can be expressed using a \forall^{k+1} sentence.

The sentence ψ_k is preserved under substructures over \mathcal{S}

We show that $\varphi_k = \neg\psi_k$ is preserved under extensions over \mathcal{S} .

Let $\mathfrak{A} \models \varphi_k$ and $\mathfrak{A} \subseteq \mathfrak{B}$. If $\alpha = (\xi_1 \wedge \xi_2 \wedge \xi_3)$ is such that $\mathfrak{A} \models \neg\alpha$, then since $\neg\alpha$ is equivalent to an existential sentence, we have $\mathfrak{B} \models \neg\alpha$; whence $\mathfrak{B} \models \varphi_k$. Else, $\mathfrak{A} \models \alpha \wedge \xi_4$. Let b be an element of \mathfrak{B} that is not in \mathfrak{A} . Then there are two possibilities:

1. $(\mathfrak{B}, a_1, b, a_2) \models ((x \leq y) \wedge (y \leq z))$ for two elements a_1, a_2 of \mathfrak{A} such that $(\mathfrak{A}, a_1, a_2) \models S(x, z)$; then $\mathfrak{B} \models \neg\xi_3$ and hence $\mathfrak{B} \models \varphi_k$.
2. $(\mathfrak{B}, b) \models ((d \leq x) \vee (x \leq c))$. Since the interpretations of c, d in \mathfrak{B} are the same as those of c, d in \mathfrak{A} respectively, we have $\mathfrak{B} \models \neg\xi_2$ and hence $\mathfrak{B} \models \varphi_k$.

In all cases, we have $\mathfrak{B} \models \varphi_k$.

The sentence ψ_k is not equivalent over \mathcal{S} to any $\exists^k\forall^*$ sentence

Suppose ψ_k is equivalent to the sentence $\gamma = \exists x_1 \dots \exists x_k \forall^n \bar{y} \beta(x_1, \dots, x_k, \bar{y})$, where β is a quantifier-free formula. We show below that this leads to a contradiction, showing that the sentence γ cannot exist.

Consider the structure $\mathfrak{A} = (\mathcal{U}_{\mathfrak{A}}, \leq^{\mathfrak{A}}, S^{\mathfrak{A}}, P^{\mathfrak{A}}, c^{\mathfrak{A}}, d^{\mathfrak{A}})$, where the universe $\mathcal{U}_{\mathfrak{A}} = \{1, \dots, (8n+1) \times (k+1)\}$, $\leq^{\mathfrak{A}}$ and $S^{\mathfrak{A}}$ are respectively the usual linear order and successor relation on $\mathcal{U}_{\mathfrak{A}}$, $c^{\mathfrak{A}} = 1, d^{\mathfrak{A}} = (8n+1) \times (k+1)$ and $P^{\mathfrak{A}} = \{(4n+1) + i \times (8n+1) \mid i \in \{0, \dots, k\}\}$. We see that $\mathfrak{A} \models (\xi_1 \wedge \xi_2 \wedge \xi_3 \wedge \xi_4 \wedge \neg\xi_5)$ and hence $\mathfrak{A} \models \psi_k$. Then $\mathfrak{A} \models \gamma$. Let a_1, \dots, a_k be the witnesses in \mathfrak{A} to the k existential quantifiers of γ .

It is clear that there exists $i^* \in \{0, \dots, k\}$ such that a_j does not belong to $\{(8n+1) \times i^* + 1, \dots, (8n+1) \times (i^* + 1)\}$ for each $j \in \{1, \dots, k\}$. Then consider the structure \mathfrak{B} that is identical to \mathfrak{A} except that $P^{\mathfrak{B}} = P^{\mathfrak{A}} \setminus \{(4n+1) + i^* \times (8n+1)\}$. It is clear from the definition of \mathfrak{B} that $\mathfrak{B} \models (\xi_1 \wedge \xi_2 \wedge \xi_3 \wedge \xi_4 \wedge \xi_5)$ and hence $\mathfrak{B} \models \neg\psi_k$. We now show a contradiction by showing that $\mathfrak{B} \models \gamma$.

We show that $\mathfrak{B} \models \gamma$ by showing that $(\mathfrak{B}, a_1, \dots, a_k) \models \forall^n \bar{y} \beta(x_1, \dots, x_k, \bar{y})$. This is in turn done by showing that for any n -tuple $\bar{e} = (e_1, \dots, e_n)$ from \mathfrak{B} , there exists an n -tuple $\bar{f} = (f_1, \dots, f_n)$ from \mathfrak{A} such that the (partial) map $\rho : \mathfrak{B} \rightarrow \mathfrak{A}$ given by $\rho(1) = 1, \rho((8n+1) \times (k+1)) = (8n+1) \times (k+1), \rho(a_j) = a_j$ for $j \in \{1, \dots, k\}$ and $\rho(e_j) = f_j$ for $j \in \{1, \dots, n\}$ is such that ρ is a partial isomorphism from \mathfrak{B} to \mathfrak{A} . Then since $(\mathfrak{A}, a_1, \dots, a_k) \models \forall^n \bar{y} \beta(x_1, \dots, x_k, \bar{y})$,

we have $(\mathfrak{A}, a_1, \dots, a_k, \bar{f}) \models \beta(x_1, \dots, x_k, \bar{y})$ whereby $(\mathfrak{B}, a_1, \dots, a_k, \bar{e}) \models \beta(x_1, \dots, x_k, \bar{y})$. Since \bar{e} is an arbitrary n -tuple from \mathfrak{B} , we have $(\mathfrak{B}, a_1, \dots, a_k) \models \forall^n \bar{y} \beta(x_1, \dots, x_k, \bar{y})$.

Define a *contiguous segment* in \mathfrak{B} to be a set of l distinct elements of \mathfrak{B} , for some $l \geq 1$, that are contiguous w.r.t. the linear ordering in \mathfrak{B} . That is, if b_1, \dots, b_l are the distinct elements of the aforesaid contiguous segment such that $(b_j, b_{j+1}) \in \leq^{\mathfrak{B}}$ for $1 \leq j \leq l-1$, then $(b_j, b_{j+1}) \in S^{\mathfrak{B}}$. We represent such a contiguous segment as $[b_1, b_l]$, and view it as an interval in \mathfrak{B} . Given an n -tuple \bar{e} from \mathfrak{B} , a *contiguous segment of \bar{e} in \mathfrak{B}* is a contiguous segment in \mathfrak{B} , all of whose elements belong to (the set underlying) \bar{e} . A *maximal contiguous segment of \bar{e} in \mathfrak{B}* is a contiguous segment of \bar{e} in \mathfrak{B} that is not strictly contained in another contiguous segment of \bar{e} in \mathfrak{B} . Let CS be the set of all maximal contiguous segments of \bar{e} in \mathfrak{B} . Let $\text{CS}_1 \subseteq \text{CS}$ be the set of all those segments of CS that have an intersection with the set $\{1, \dots, (8n+1) \times i^*\} \cup \{(8n+1) \times (i^*+1) + 1, \dots, (8n+1) \times (k+1)\}$. Let $\text{CS}_2 = \text{CS} \setminus \text{CS}_1$. Then all intervals in CS_2 are strictly contained in the interval $[(8n+1) \times i^* + 1, (8n+1) \times (i^*+1)]$. Let $\text{CS}_2 = \{[i_1, j_1], [i_2, j_2], \dots, [i_r, j_r]\}$. Observe that $r \leq n$. Without loss of generality, assume that $i_1 \leq j_1 < i_2 \leq j_2 < \dots < i_r \leq j_r$. Let CS_3 be the set of contiguous segments in \mathfrak{A} defined as $\text{CS}_3 = \{[i'_1, j'_1], [i'_2, j'_2], \dots, [i'_r, j'_r]\}$ where $i'_1 = (8n+1) \times i^* + n + 1$, $j'_1 = i'_1 + (j_1 - i_1)$, and for $2 \leq l \leq r$, we have $i'_l = j'_{l-1} + 2$ and $j'_l = i'_l + (j_l - i_l)$. Observe that the sum of the lengths of the segments of CS_2 is at most n , whereby $j'_r \leq (8n+1) \times i^* + 3n + 1$.

Now consider the tuple $\bar{f} = (f_1, \dots, f_n)$ defined using $\bar{e} = (e_1, \dots, e_n)$ as follows. Let $\text{Elements}(\text{CS}_1)$, resp. $\text{Elements}(\text{CS}_2)$, denote the elements contained in the segments of CS_1 , resp. CS_2 . For $1 \leq l \leq n$, if $e_l \in \text{Elements}(\text{CS}_1)$, then $f_l = e_l$. Else suppose e_l belongs to the segment $[i_s, j_s]$ of CS_2 where $1 \leq s \leq r$, and suppose that $e_l = i_s + t$ for some $t \in \{0, (j_s - i_s)\}$. Then choose $f_l = i'_s + t$.

It is easy to see that the (partial) map $\rho : \mathfrak{B} \rightarrow \mathfrak{A}$ given by $\rho(1) = 1$, $\rho((8n+1) \times (k+1)) = (8n+1) \times (k+1)$, $\rho(a_j) = a_j$ for $j \in \{1, \dots, k\}$ and $\rho(e_j) = f_j$ for $j \in \{1, \dots, n\}$ is such that ρ is a partial isomorphism from \mathfrak{B} to \mathfrak{A} . \square

Remark 8.1.2. Proposition 8.1.1 is a stronger statement than the failure of the Łoś-Tarski theorem in the finite. While the latter only shows that the class of sentences that are preserved under substructures in the finite, i.e. the class of sentences that are $PSC(0)$ in the finite, cannot be characterized by the class of \forall^* sentences, Proposition 8.1.1 shows that for each $l \geq 0$, the class of sentences that is $PSC(l)$ in the finite cannot be characterized by, or even semantically subsumed by, the class of $\exists^k \forall^*$ sentences, for any $k \geq 0$. Interestingly, the sentence ψ_k in

Proposition 8.1.1, which is not equivalent in the finite to any $\exists^k \forall^*$ sentence, is actually equivalent (in the finite) to an $\exists^{k+1} \forall^*$ sentence. We dwell on this observation towards the end of Chapter 12.

8.2 Failure of $\text{GLT}(k)$ over classes that are well-behaved w.r.t. the Łoś-Tarski theorem

Towards the central result of this section, we first show the following.

Lemma 8.2.1. *Let \mathcal{U} be the class of all undirected graphs that are (finite) disjoint unions of (finite) undirected paths. Let \mathcal{S} be a class of undirected graphs of degree at most 2, that contains \mathcal{U} as a subclass. Then for each $k \geq 2$, there is a sentence ϕ_k that is $\text{PSC}(k)$ over \mathcal{S} , but that is not \mathcal{S} -equivalent to any $\exists^k \forall^*$ sentence.*

Proof. Given $k \geq 2$, consider ϕ_k that asserts that either (i) there are at least k nodes of degree exactly 0 or (ii) there are at least $k + 1$ nodes of degree at most 1. We claim that the sentence ϕ_k is the desired sentence for the given k . We give the reasoning for the case of $k = 2$; an analogous reasoning can be done for $k > 2$. In our arguments below, $\phi = \phi_2$. We observe that any graph in \mathcal{S} is a disjoint union of undirected paths and undirected cycles.

Any graph in \mathcal{S} that contains a single connected component that is a path (whereby every other connected component is a cycle), cannot be a model of ϕ . Then every model \mathfrak{D} of ϕ in \mathcal{S} has at least two connected components, each of which is a path (of length ≥ 0). Consider a set C formed by an end point of one of these paths and an end point of the other of these paths. It is easy to check that C is a 2-crux of \mathfrak{D} w.r.t. ϕ over \mathcal{S} , whereby ϕ is $\text{PSC}(2)$ over \mathcal{S} . Suppose ϕ is \mathcal{S} -equivalent to $\psi = \exists x_1 \exists x_2 \forall^n \bar{y} \beta(x_1, x_2, \bar{y})$ where β is a quantifier-free formula. Consider a graph $\mathfrak{A} \in \mathcal{U}$ that has exactly two connected components each of which is a path of length $\geq 2n$. Clearly \mathfrak{A} is a model of ϕ . Further since $\mathfrak{A} \in \mathcal{S}$, we have $\mathfrak{A} \models \psi$. Let a_1, a_2 be witnesses in \mathfrak{A} , to the existential quantifiers of ψ . It cannot be that a_1, a_2 are both from the same path of \mathfrak{A} else the path by itself would be a model for ψ , and hence ϕ . Now consider a structure $\mathfrak{B} \in \mathcal{U} \subseteq \mathcal{S}$, which is a single path that has length $\geq 4n$, and let p_1, p_2 be the end points of this path. If a_1 (resp. a_2) is at a distance of $r \leq n$ from the end point of any path in \mathfrak{A} , then choose a point b_1 (resp. b_2) at the same distance, namely r , from p_1 (resp. p_2) in \mathfrak{B} . Else choose b_1 (resp. b_2) at a distance of $n + 1$ from p_1 (resp. p_2). Now consider any n -tuple \bar{e} from \mathfrak{B} . By a similar kind

of reasoning as done in Proposition 8.1.1, one can show that it is possible to choose an n -tuple \bar{f} from \mathfrak{A} such that the (partial) map $\rho : \mathfrak{B} \rightarrow \mathfrak{A}$ given by $\rho(b_i) = a_i$, $\rho(e_j) = f_j$ for $1 \leq i \leq 2$ and $1 \leq j \leq n$, is a partial isomorphism from \mathfrak{B} to \mathfrak{A} . Since $(\mathfrak{A}, a_1, a_2) \models \forall^n \bar{y} \beta(x_1, x_2, \bar{y})$, we have $(\mathfrak{A}, a_1, a_2, \bar{f}) \models \beta(x_1, x_2, \bar{y})$ whereby $(\mathfrak{B}, b_1, b_2, \bar{e}) \models \beta(x_1, x_2, \bar{y})$. Since \bar{e} is arbitrary, we have \mathfrak{B} models ψ , and hence ϕ – a contradiction. \square

To state the central result of this section, we first introduce some terminology. Let τ be a vocabulary consisting of unary and binary relation symbols only. Given a τ -structure \mathfrak{A} , let $\mathcal{G}(\mathfrak{A}) = (V, E)$ be the graph such that (i) V is exactly the universe of \mathfrak{A} , and (ii) $(a, b) \in E$ iff for some binary relation symbol $B \in \tau$, we have $(\mathfrak{A}, a, b) \models (B(x, y) \vee B(y, x))$. In the language of translation schemes (see Section 7.4), the structure $\mathcal{G}(\mathfrak{A})$ is indeed $\Xi(\mathfrak{A})$ where $\Xi = (\xi, \xi_E)$ is the $(1, \tau, \{E\}, \text{FO})$ -translation scheme such that ξ is the formula $x = x$, and ξ_E is the formula $\bigvee_{D \in \tau_{bin}} (D(x, y) \vee D(y, x))$, where τ_{bin} is the set of all binary relation symbols of τ . Given a class \mathcal{S} of τ -structures, let $\mathcal{G}(\mathcal{S}) = \{\mathcal{G}(\mathfrak{A}) \mid \mathfrak{A} \in \mathcal{S}\}$.

We say a class of undirected graphs has *unbounded induced path lengths* if for every $n \in \mathbb{N}$, there exists a graph G in the class such that G contains an induced path of length $\geq n$. We say a class \mathcal{S} of τ -structures has unbounded induced path lengths if the class $\mathcal{G}(\mathcal{S})$ of undirected graphs has unbounded induced path lengths. A class of τ -structures is said to have *bounded induced path lengths* if it does not have unbounded induced path lengths.

The central result of this section is as follows.

Theorem 8.2.2. *Let \mathcal{V} be a hereditary class of undirected graphs. Let τ be a vocabulary containing unary and binary relation symbols only, and \mathcal{S} be the class of all τ -structures \mathfrak{A} such that $\mathcal{G}(\mathfrak{A})$ belongs to \mathcal{V} . If $\text{GLT}(k)$ holds over \mathcal{S} for any $k \geq 2$, then \mathcal{S} has bounded induced path lengths.*

Proof. If τ contains only unary relation symbols, then trivially \mathcal{S} has bounded induced path lengths; the induced path lengths in all structures in \mathcal{S} is bounded by 0. Therefore, assume τ contains at least one binary relation symbol. Let τ_{bin} be the set of all binary relation symbols of τ , and let B be one such relation symbol of τ_{bin} .

We prove the result by contradiction. Suppose \mathcal{S} has unbounded induced path lengths. Then for every $n \in \mathbb{N}$, there exists $\mathfrak{A} \in \mathcal{S}$ such that the graph $\mathcal{G}(\mathfrak{A})$ contains an induced path of length $r \geq n$. Since $\mathcal{G}(\mathfrak{A}) \in \mathcal{V}$ and \mathcal{V} is hereditary, it follows that the undirected path graph of length n belongs to \mathcal{V} , for each $n \in \mathbb{N}$. Then, again by the hereditariness of \mathcal{V} , the class \mathcal{U} of all (finite)

disjoint unions of (finite) undirected paths is a subclass of \mathcal{V} . Let χ be a universal sentence in the vocabulary $\{E\}$ of graphs such that any model of χ is an undirected graph whose degree is at most 2, and let \mathcal{V}_1 be the class of all models of χ in \mathcal{V} . Clearly, \mathcal{U} is a subclass of \mathcal{V}_1 . For $k \geq 2$, let ϕ_k be the sentence given by Lemma 8.2.1 such that ϕ_k is $PSC(k)$ over \mathcal{V}_1 but ϕ_k is not \mathcal{V}_1 -equivalent to any $\exists^k\forall^*$ sentence.

Before we proceed, we present two observations. Let $\xi_E(x, y) = \bigvee_{D \in \tau_{bin}} (D(x, y) \vee D(y, x))$ as seen earlier. Given an $\text{FO}(\{E\})$ sentence β , let $\beta[E \mapsto \xi_E]$ be the $\text{FO}(\tau)$ sentence obtained from β by replacing each occurrence of “ $E(x, y)$ ” in β , with the formula $\xi_E(x, y)$. Following are two observations that are easy to verify. Let $\mathfrak{A}, \mathfrak{B}$ be given τ -structures.

O.1 If $\mathfrak{B} \subseteq \mathfrak{A}$, then $\mathcal{G}(\mathfrak{B}) \subseteq \mathcal{G}(\mathfrak{A})$.

O.2 $\mathfrak{A} \models \beta[E \mapsto \xi_E]$ iff $\mathcal{G}(\mathfrak{A}) \models \beta$.

(Indeed, Observation O.2 can also be verified using Proposition 7.4.2, while Observation O.1 can also be verified using Lemma 10.4.10, that we present later.)

Consider now the $\text{FO}(\tau)$ sentence $\alpha = (\phi_k \wedge \chi)[E \mapsto \xi_E]$. We have the following.

α is $PSC(k)$ over \mathcal{S} : Suppose $\mathfrak{A} \in \mathcal{S}$ is such that $\mathfrak{A} \models \alpha$. Then $\mathcal{G}(\mathfrak{A}) \in \mathcal{V}$ and $\mathcal{G}(\mathfrak{A}) \models \phi_k \wedge \chi$ by Observation O.2; whereby $\mathcal{G}(\mathfrak{A}) \in \mathcal{V}_1$. Then since ϕ_k is $PSC(k)$ over \mathcal{V}_1 , there exists a k -crux C of $\mathcal{G}(\mathfrak{A})$ w.r.t. ϕ_k over \mathcal{V}_1 . Consider a substructure \mathfrak{B} of \mathfrak{A} such that $\mathfrak{B} \in \mathcal{S}$ and \mathfrak{B} contains C . Then $\mathcal{G}(\mathfrak{B}) \in \mathcal{V}$. By Observation O.1 above, we have $\mathcal{G}(\mathfrak{B}) \subseteq \mathcal{G}(\mathfrak{A})$. Since χ is a universal sentence that is true in $\mathcal{G}(\mathfrak{A})$, we have that $\mathcal{G}(\mathfrak{B})$ models χ , whereby $\mathcal{G}(\mathfrak{B}) \in \mathcal{V}_1$. Then since $\mathcal{G}(\mathfrak{B})$ contains C and C is a k -crux of $\mathcal{G}(\mathfrak{A})$ w.r.t. ϕ_k over \mathcal{V}_1 , it follows that $\mathcal{G}(\mathfrak{B}) \models (\phi_k \wedge \chi)$. By Observation O.2 above, $\mathfrak{B} \models \alpha$. Whereby C is a k -crux of \mathfrak{A} w.r.t. α over \mathcal{S} . Then α is $PSC(k)$ over \mathcal{S} .

α is not \mathcal{S} -equivalent to any $\exists^k\forall^*$ sentence: Suppose α is \mathcal{S} -equivalent to an $\exists^k\forall^*$ $\text{FO}(\tau)$ sentence γ_1 . Let $\gamma_2 = \gamma_1[B \mapsto E; D \mapsto \text{False}, D \in \tau_{bin}, D \neq B]$ be the $\text{FO}(\{E\})$ sentence obtained from γ_1 by replacing (i) each occurrence of B in γ_1 with E , and (ii) each occurrence of D in γ_1 with the constant formula False , for each $D \in \tau_{bin}, D \neq B$. Observe that γ_2 is an $\exists^k\forall^*$ sentence. We show that ϕ_k is \mathcal{V}_1 -equivalent to γ_2 , contradicting Lemma 8.2.1.

Let ψ be the $\text{FO}(\tau)$ sentence given by $\psi = \forall x \forall y ((B(x, y) \leftrightarrow B(y, x)) \wedge \bigwedge_{D \in \tau_{bin}, D \neq B} \neg D(x, y))$. Given a $\{E\}$ -structure \mathfrak{C} , let \mathfrak{C}^ψ be the τ -structure such that (i) $\mathcal{U}_{\mathfrak{C}^\psi} = \mathcal{U}_{\mathfrak{C}}$, (ii) for any two elements $a, b \in \mathcal{U}_{\mathfrak{C}^\psi}$, $(\mathfrak{C}^\psi, a, b) \models B(x, y)$ iff $(\mathfrak{C}, a, b) \models E(x, y)$, and (iii) for any two elements $a, b \in \mathcal{U}_{\mathfrak{C}^\psi}$, $(\mathfrak{C}^\psi, a, b) \models \neg D(x, y)$ for each $D \in \tau_{bin}, D \neq B$. It is easy to see that $\mathfrak{C}^\psi \models \psi$, and that $\mathcal{G}(\mathfrak{C}^\psi) = \mathfrak{C}$.

Before proceeding to show that ϕ_k is \mathcal{V}_1 -equivalent to γ_2 , we make the simple yet important observation, call it (\dagger) , that any model of ψ in \mathcal{S} models the sentence $(\gamma_1 \leftrightarrow \gamma_2 [E \mapsto \xi_E])$.

- ϕ_k entails γ_2 over \mathcal{V}_1 : Suppose $\mathfrak{C} \in \mathcal{V}_1$ is such that $\mathfrak{C} \models \phi_k$. Then $\mathfrak{C} \models (\phi_k \wedge \chi)$. Let $\mathfrak{A} = \mathfrak{C}^\psi$. Then $\mathfrak{A} \models \psi$ and $\mathcal{G}(\mathfrak{A}) = \mathfrak{C}$, whereby $\mathfrak{A} \in \mathcal{S}$. Since $\mathfrak{C} \models (\phi_k \wedge \chi)$, we have $\mathfrak{A} \models \alpha$ by Observation O.2 above. Now since α is \mathcal{S} -equivalent to γ_1 (by assumption), we have $\mathfrak{A} \models \gamma_1$. Since \mathfrak{A} models ψ , we have by (\dagger) that $\mathfrak{A} \models \gamma_2 [E \mapsto \xi_E]$. By Observation O.2 above, $\mathfrak{C} \models \gamma_2$.
- γ_2 entails ϕ_k over \mathcal{V}_1 : Suppose $\mathfrak{C} \in \mathcal{V}_1$ is such that $\mathfrak{C} \models \gamma_2$. Let $\mathfrak{A} = \mathfrak{C}^\psi$. Then $\mathfrak{A} \models \psi$ and $\mathcal{G}(\mathfrak{A}) = \mathfrak{C}$. By Observation O.2 above, $\mathfrak{A} \models \gamma_2 [E \mapsto \xi_E]$. By (\dagger) , we have $\mathfrak{A} \models \gamma_1$. Since γ_1 is \mathcal{S} -equivalent to α , we have $\mathfrak{A} \models \alpha$. By Observation O.2, $\mathfrak{C} \models (\phi_k \wedge \chi)$, and hence $\mathfrak{C} \models \phi_k$.

□

Proposition 8.2.3. *Let τ be a vocabulary containing unary and binary relation symbols only, and let there be at least one binary relation symbol in τ . Then there exist classes \mathcal{S}_1 and \mathcal{S}_2 of τ -structures such that,*

- \mathcal{S}_1 is acyclic, of degree at most 2, and is closed under substructures and disjoint unions
- \mathcal{S}_2 is the class of all τ -structures of treewidth 1

and $\text{GLT}(k)$ fails over each of \mathcal{S}_1 and \mathcal{S}_2 for each $k \geq 2$.

Proof. Let \mathcal{V}_2 be the class of all undirected graphs that are acyclic, and let \mathcal{V}_1 be the class of all the graphs in \mathcal{V}_1 that have degree at most 2. Clearly \mathcal{V}_1 and \mathcal{V}_2 are hereditary. Let \mathcal{S}_i be the class of all τ -structures \mathfrak{A} such that $\mathcal{G}(\mathfrak{A}) \in \mathcal{V}_i$, for $i \in \{1, 2\}$. It is easy to see that \mathcal{S}_1 is acyclic, of degree at most 2, and is closed under substructures and disjoint unions. That \mathcal{S}_2 is the class of all τ -structures of tree-width 1 follows from definitions. Observe now that each of \mathcal{S}_1 and \mathcal{S}_2 has unbounded induced path lengths. It then follows from Theorem 8.2.2, that $\text{GLT}(k)$ cannot hold over \mathcal{S}_i for any $k \geq 2$, for each $i \in \{1, 2\}$. □

The above result motivates us to ask the following question: *Can we identify structural properties (possibly abstract) of classes of finite structures that are satisfied by interesting classes of finite structures, and that entail $\text{GLT}(k)$? And further, entail $\text{GLT}(k)$ in effective form?.* We identify one such property in this thesis. Notably, the classes of structures that satisfy our property are incomparable to those studied in [7, 21, 38].

Chapter 9

The \mathcal{L} -Equivalent Bounded Substructure Property – \mathcal{L} -EBSP(\mathcal{S}, k)

We define a new logic-based combinatorial property of classes of finite structures.

Definition 9.1 (\mathcal{L} -EBSP(\mathcal{S}, k)). Let \mathcal{S} be a class of structures and $k \in \mathbb{N}$. We say that \mathcal{S} satisfies the \mathcal{L} -equivalent bounded substructure property for parameter k , abbreviated \mathcal{L} -EBSP(\mathcal{S}, k) is true (alternatively, \mathcal{L} -EBSP(\mathcal{S}, k) holds), if there exists a function $\theta_{(\mathcal{S}, k, \mathcal{L})} : \mathbb{N} \rightarrow \mathbb{N}$ such that for each $m \in \mathbb{N}$, for each structure \mathfrak{A} of \mathcal{S} and for each k -tuple \bar{a} from \mathfrak{A} , there exists a structure \mathfrak{B} such that (i) $\mathfrak{B} \in \mathcal{S}$, (ii) $\mathfrak{B} \subseteq \mathfrak{A}$, (iii) the elements of \bar{a} are contained in \mathfrak{B} , (iv) $|\mathfrak{B}| \leq \theta_{(\mathcal{S}, k, \mathcal{L})}(m)$, and (v) $\text{tp}_{\mathfrak{B}, \bar{a}, m, \mathcal{L}}(\bar{x}) = \text{tp}_{\mathfrak{A}, \bar{a}, m, \mathcal{L}}(\bar{x})$. The conjunction of these five conditions is denoted as \mathcal{L} -EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, k, m, \bar{a}, \theta_{(\mathcal{S}, k, \mathcal{L})}$). We call $\theta_{(\mathcal{S}, k, \mathcal{L})}$ a *witness function* for \mathcal{L} -EBSP(\mathcal{S}, k).

Remark 9.2. Observe that there can be several witness functions for \mathcal{L} -EBSP(\mathcal{S}, k). For instance, if $\theta_{(\mathcal{S}, k, \mathcal{L})}$ is a witness function, then any function $\nu : \mathbb{N} \rightarrow \mathbb{N}$ such that $\theta_{(\mathcal{S}, k, \mathcal{L})}(m) \leq \nu(m)$ for all $m \in \mathbb{N}$ is also a witness function. Observe also that there always exists a monotonic witness function. This is easily seen as follows. For a function $\nu : \mathbb{N} \rightarrow \mathbb{N}$, let $\nu' : \mathbb{N} \rightarrow \mathbb{N}$ be the function defined as $\nu'(m) = \sum_{i=0}^m \nu(i)$. Then ν' is monotonic and $\nu(m) \leq \nu'(m)$ for all $m \in \mathbb{N}$; whereby if ν is a witness function, then so is ν' . Therefore, we assume henceforth that all witness functions are monotonic.

We list below two simple examples of classes \mathcal{S} that satisfy \mathcal{L} -EBSP(\mathcal{S}, k) for every $k \in \mathbb{N}$. Many more such classes are presented in Chapter 10.

1. Let \mathcal{S} be a finite class of structures. Clearly, \mathcal{L} -EBSP(\mathcal{S}, k) holds for all $k \in \mathbb{N}$, with $\theta_{(\mathcal{S}, k, \mathcal{L})}(m)$ giving the size of the largest structure in \mathcal{S} .
2. Let \mathcal{S} be the class of all τ -structures, where all relation symbols in τ are unary. By a simple FO-EF game argument, one can see that FO-EBSP(\mathcal{S}, k) holds for all $k \in \mathbb{N}$, with $\theta_{(\mathcal{S}, k, \mathcal{L})}(m) = m \cdot 2^{|\tau|} + k$. In more detail: given $\mathfrak{A} \in \mathcal{S}$, one can associate exactly one of $2^{|\tau|}$ colors with each element a of \mathfrak{A} , where the colour of a gives the valuation in \mathfrak{A} , of all predicates of τ for a . Then given a k -tuple \bar{a} from \mathfrak{A} , consider $\mathfrak{B} \subseteq \mathfrak{A}$ satisfying the following: (i) the set W of elements of \bar{a} , is contained in $U_{\mathfrak{B}}$, and (ii) for each colour c , if $A_c = \{a \mid a \in U_{\mathfrak{A}} \setminus W, a \text{ has colour } c \text{ in } \mathfrak{A}\}$, then $A_c \subseteq U_{\mathfrak{B}}$ if $|A_c| < m$, else $|A_c \cap U_{\mathfrak{B}}| = m$. It is then easy to see that FO-EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, k, m, \bar{a}, \theta_{(\mathcal{S}, k, \mathcal{L})}$) is true.

By a similar MSO-EF game argument, one can show that MSO-EBSP(\mathcal{S}, k) holds for all $k \in \mathbb{N}$, with the same witness function $\theta_{(\mathcal{S}, k, \mathcal{L})}$ as above.

9.1 \mathcal{L} -EBSP(\cdot, k) entails GLT(k)

In this section, we show that \mathcal{L} -EBSP(\mathcal{S}, k) indeed entails GLT(k). Towards this result, we first observe that given a class \mathcal{S} of structures and a natural number n , there exists an FO sentence that defines the subclass of all structures in \mathcal{S} of size at most n (it is easy to construct such a sentence that is in $\exists^n \forall^*$). We fix such a sentence and denote it as $\xi_{\mathcal{S}, n}$. Secondly, we observe the following.

Lemma 9.1.1. *Given a sentence ψ over a vocabulary τ and variables $\bar{x} = (x_1, \dots, x_n)$ not appearing in ψ , there exists a quantifier-free formula $\psi|_{\bar{x}}(\bar{x})$ over τ , whose free variables are among \bar{x} , such that the following holds: Let \mathfrak{A} be a structure and $\bar{a} = (a_1, \dots, a_n)$ a sequence of elements of \mathfrak{A} . Then*

$$(\mathfrak{A}, a_1, \dots, a_n) \models \psi|_{\bar{x}}(\bar{x}) \text{ iff } \mathfrak{A}(\{a_1, \dots, a_k\}) \models \psi$$

where $\mathfrak{A}(\{a_1, \dots, a_k\})$ denotes the substructure of \mathfrak{A} induced by $\{a_1, \dots, a_k\}$. Further, $\psi|_{\bar{x}}(\bar{x})$ is computable from ψ .

Proof. Let $X = \{x_1, \dots, x_n\}$. Replace every subformula of ψ of the form $\exists z \chi(z, y_1, \dots, y_k)$ with $\bigvee_{z \in X} \chi(z, y_1, \dots, y_k)$, and every subformula of ψ of the form $\forall z \chi(z, y_1, \dots, y_k)$ with

$\bigwedge_{z \in X} \chi(z, y_1, \dots, y_k)$. It is clear that the resulting formula can be taken to be $\psi|_{\bar{x}}(\bar{x})$. It is also clear that $\psi|_{\bar{x}}(\bar{x})$ is computable from ψ . \square

The formula $\psi|_{\bar{x}}(\bar{x})$ is read as ψ *relativized to* \bar{x} .

Given a class \mathcal{S} of structures and an \mathcal{L} sentence φ , we say that φ is $PSC(k)$ over \mathcal{S} if the class of models of φ in \mathcal{S} is $PSC(k)$ over \mathcal{S} (see Definition 3.1.1 for the definition of $PSC(k)$). We say \mathcal{L} -GLT(k) holds over \mathcal{S} if for all \mathcal{L} sentences φ , it is the case that φ is $PSC(k)$ over \mathcal{S} iff φ is \mathcal{S} -equivalent to an $\exists^k \forall^*$ FO sentence. Observe that over any class \mathcal{S} , FO-GLT(k) holds iff GLT(k) holds, and that if MSO-GLT(k) holds, then so does FO-GLT(k).

Theorem 9.1.2. *Let \mathcal{S} be a class of finite structures and $k \in \mathbb{N}$ be such that \mathcal{L} -EBSP(\mathcal{S}, k) holds. Then \mathcal{L} -GLT(k), and hence GLT(k), holds over \mathcal{S} . Further, if there exists a computable witness function for \mathcal{L} -EBSP(\mathcal{S}, k), then the translation from an \mathcal{L} sentence that is $PSC(k)$ over \mathcal{S} , to an \mathcal{S} -equivalent $\exists^k \forall^*$ sentence, is effective.*

Proof. Let $\theta_{(\mathcal{S}, k, \mathcal{L})}$ be a witness function for \mathcal{L} -EBSP(\mathcal{S}, k). We show below that an \mathcal{L} sentence φ of quantifier rank m that is $PSC(k)$ over \mathcal{S} , is \mathcal{S} -equivalent to the sentence χ given by $\chi = \exists^k \bar{x} \forall^p \bar{y} \psi|_{\bar{x}\bar{y}}(\bar{x}, \bar{y})$, where $p = \theta_{(\mathcal{S}, k, \mathcal{L})}(m)$ and $\psi = (\xi_{\mathcal{S}, p} \rightarrow \varphi)$. It is easy to see that if $\theta_{(\mathcal{S}, k, \mathcal{L})}$ is computable, then since m is effectively computable from φ , so are p , $\xi_{\mathcal{S}, p}$ and χ .

From the discussion in Section 3.1, we see that a \mathcal{L} sentence that is \mathcal{S} -equivalent to an $\exists^k \forall^*$ sentence, is $PSC(k)$ over \mathcal{S} . Towards the converse, consider an \mathcal{L} sentence φ , of quantifier rank m , that is $PSC(k)$ over \mathcal{S} . Let χ be the sentence as given in the previous paragraph. Since φ is $PSC(k)$ over \mathcal{S} , every model \mathfrak{A} of φ in \mathcal{S} also satisfies χ . This is because the elements of any k -crux of \mathfrak{A} can serve as witnesses to the existential quantifiers of χ . Thus φ entails χ over \mathcal{S} . To show χ entails φ over \mathcal{S} , suppose \mathfrak{A} is a model of χ in \mathcal{S} . Let \bar{a} be a k -tuple that is a witness in \mathfrak{A} to the k existential quantifiers of χ . Since \mathcal{L} -EBSP(\mathcal{S}, k) holds, there exists a structure \mathfrak{B} such that \mathcal{L} -EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, k, m, \bar{a}, \theta_{(\mathcal{S}, k, \mathcal{L})}$) is true. In other words, we have (i) $\mathfrak{B} \in \mathcal{S}$, (ii) $\mathfrak{B} \subseteq \mathfrak{A}$, (iii) the elements of \bar{a} are contained in \mathfrak{B} , (iv) $|\mathfrak{B}| \leq \theta_{(\mathcal{S}, k, \mathcal{L})}(m) = p$, and (v) $(\mathfrak{B}, \bar{a}) \equiv_{m, \mathcal{L}} (\mathfrak{A}, \bar{a})$; then $\mathfrak{B} \equiv_{m, \mathcal{L}} \mathfrak{A}$. Since $(\mathfrak{A}, \bar{a}) \models \forall^p \bar{y} \psi|_{\bar{x}\bar{y}}(\bar{x}, \bar{y})$, by instantiating the universal variables \bar{y} with the elements of $U_{\mathfrak{B}}$, and by using Lemma 9.1.1, we get $\mathfrak{B} \models (\xi_{\mathcal{S}, p} \rightarrow \varphi)$. Since $\mathfrak{B} \in \mathcal{S}$ and $|\mathfrak{B}| \leq p$, we have $\mathfrak{B} \models \xi_{\mathcal{S}, p}$ whereby $\mathfrak{B} \models \varphi$. Since φ is an \mathcal{L} sentence of quantifier rank m and $\mathfrak{B} \equiv_{m, \mathcal{L}} \mathfrak{A}$, it follows that $\mathfrak{A} \models \varphi$. Thus χ entails φ over \mathcal{S} whereby φ is \mathcal{S} -equivalent to χ . \square

9.2 \mathcal{L} -EBSP(\mathcal{S}, k) – a finitary analogue of the downward

Löwenheim-Skolem property

The downward Löwenheim-Skolem theorem, as already seen in Section 2.3, is one of the first results of classical model theory. Well before the central tool of classical model theory, namely the compactness theorem, was discovered, Löwenheim and Skolem [56, 78], showed the following result.

Theorem 9.2.1 (Löwenheim 1915, Skolem 1920). *If an FO theory over a countable vocabulary has an infinite model, then it has a countable model.*

Historically, the proof of Theorem 9.2.1, initially due to Löwenheim, assumed König’s lemma, though the latter lemma was proven only in 1927. Skolem in 1920 gave the first fully self-contained proof of the theorem and hence the theorem is jointly attributed to Skolem [8]. In subsequent years, Skolem came up with a more general statement.

Theorem 9.2.2 (Skolem’s revised version of the downward Löwenheim-Skolem theorem). *Let τ be a countable vocabulary. For every τ -structure \mathfrak{A} and every countable set W of elements of \mathfrak{A} , there is a countable substructure \mathfrak{B} of \mathfrak{A} such that \mathfrak{B} contains the elements of W , and \mathfrak{B} is elementarily equivalent to \mathfrak{A} .*

Finally, Mal’tsev [58] proved the most general version of downward Löwenheim-Skolem theorem, which is also considered as the modern statement of the theorem. This version by Mal’tsev, stated as Theorem 2.3.3 in Section 2.3, is restated below. Recall, that $\mathfrak{B} \preceq \mathfrak{A}$ denotes that \mathfrak{B} is an elementary substructure of \mathfrak{A} .

Theorem 9.2.3 (Modern statement of the Löwenheim-Skolem theorem, Mal’tsev 1936). *Let τ be a countable vocabulary. For every τ -structure \mathfrak{A} and every infinite cardinal κ , if W is a set of at most κ elements of \mathfrak{A} , then there exists a structure \mathfrak{B} such that (i) $\mathfrak{B} \subseteq \mathfrak{A}$, (ii) \mathfrak{B} contains the elements of W , (iii) $|\mathfrak{B}| \leq \kappa$, and (iv) $\mathfrak{B} \preceq \mathfrak{A}$.*

We now define a model-theoretic property of arbitrary structures, that is closely related to the model-theoretic properties contained in the versions of the downward Löwenheim-Skolem theorem by Skolem and Mal’tsev. Given structures \mathfrak{A} and \mathfrak{B} , we say \mathfrak{A} and \mathfrak{B} are \mathcal{L} -equivalent, denoted $\mathfrak{A} \equiv_{\mathcal{L}} \mathfrak{B}$, if \mathfrak{A} and \mathfrak{B} agree on all \mathcal{L} sentences. If $\mathfrak{B} \subseteq \mathfrak{A}$ and \bar{b} is a (potentially infinite) tuple from \mathfrak{B} , that contains exactly the elements of \mathfrak{B} , then we say \mathfrak{B} is an \mathcal{L} -substructure of

\mathfrak{A} , denoted $\mathfrak{B} \preceq_{\mathcal{L}} \mathfrak{A}$, if $(\mathfrak{B}, \bar{b}) \equiv_{\mathcal{L}} (\mathfrak{A}, \bar{b})$. The reader can recognize that when $\mathcal{L} = \text{FO}$, then $\equiv_{\mathcal{L}}$ and $\preceq_{\mathcal{L}}$ are exactly the literature notions of elementary equivalence and elementary substructure (see Section 2.3). One easily sees that if $\mathfrak{B} \subseteq \mathfrak{A}$ and \bar{a} is any tuple (of any length) from \mathfrak{B} , then

$$\mathfrak{B} \preceq_{\mathcal{L}} \mathfrak{A} \rightarrow (\mathfrak{B}, \bar{a}) \equiv_{\mathcal{L}} (\mathfrak{A}, \bar{a}) \rightarrow \mathfrak{B} \equiv_{\mathcal{L}} \mathfrak{A}$$

where \rightarrow denotes the usual implication. Consider now the following model-theoretic property of a class \mathcal{S} of *arbitrary* structures. Below, a κ -tuple is a tuple of length κ .

Definition 9.2.4. Let \mathcal{S} be a class of arbitrary structures over a countable vocabulary, and let κ be an infinite cardinal. We say that $\mathcal{L}\text{-DLSP}(\mathcal{S}, \kappa)$ is *true*, if for each structure $\mathfrak{A} \in \mathcal{S}$ and each κ -tuple \bar{a} from \mathfrak{A} , there exists a structure \mathfrak{B} such that (i) $\mathfrak{B} \in \mathcal{S}$, (ii) $\mathfrak{B} \subseteq \mathfrak{A}$, (iii) the elements of \bar{a} are contained in \mathfrak{B} , (iv) $|\mathfrak{B}| \leq \kappa$, and (v) $(\mathfrak{B}, \bar{a}) \equiv_{\mathcal{L}} (\mathfrak{A}, \bar{a})$.

Let $\mathcal{L}\text{-DLSP}_{\mathcal{M}}(\mathcal{S}, \kappa)$, resp. $\mathcal{L}\text{-DLSP}_{\mathcal{S}}(\mathcal{S}, \kappa)$, be the properties obtained from Definition 9.2.4 by simply replacing the last condition in the definition with “ $\mathfrak{B} \preceq_{\mathcal{L}} \mathfrak{A}$ ”, resp. “ $\mathfrak{B} \equiv_{\mathcal{L}} \mathfrak{A}$ ”. The implication above then shows that

$$\mathcal{L}\text{-DLSP}_{\mathcal{M}}(\mathcal{S}, \kappa) \rightarrow \mathcal{L}\text{-DLSP}(\mathcal{S}, \kappa) \rightarrow \mathcal{L}\text{-DLSP}_{\mathcal{S}}(\mathcal{S}, \kappa)$$

Observe now that by taking \mathcal{L} as FO and \mathcal{S} as the class of all (i.e. finite and infinite) structures, both $\mathcal{L}\text{-DLSP}_{\mathcal{M}}(\mathcal{S}, \kappa)$ and $\mathcal{L}\text{-DLSP}_{\mathcal{S}}(\mathcal{S}, \omega)$ are true, since indeed, these are respectively, the versions of the downward Löwenheim-Skolem theorem by Mal'tsev and Skolem, given by Theorem 9.2.3 and Theorem 9.2.2. Whereby, we can see $\mathcal{L}\text{-DLSP}(\mathcal{S}, \kappa)$ as a version of the downward Löwenheim-Skolem property that is “intermediate” between the versions of the downward Löwenheim-Skolem property by Mal'tsev and Skolem. And now, as the figure below shows, $\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ reads very much like $\mathcal{L}\text{-DLSP}(\mathcal{S}, \kappa)$. Indeed then, $\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ can very well be regarded as a finitary analogue of the downward Löwenheim-Skolem property.




Figure 9.1: \mathcal{L} -EBSP(\mathcal{S}, k) as a finitary analogue of \mathcal{L} -DLSP(\mathcal{S}, κ)

9.3 A sufficient condition for \mathcal{L} -EBSP(\cdot, k)

For a class \mathcal{S} of τ -structures, let $\mathcal{S}_p = \{(\mathfrak{A}, \nu) \mid \mathfrak{A} \in \mathcal{S}, \nu : \cup \mathfrak{A} \rightarrow \{0, \dots, p-1\}\}$ be the class of all structures obtained by labeling the elements of the structures of \mathcal{S} , with elements from $\{0, \dots, p-1\}$. Formally, each structure from \mathcal{S}_p can be seen as a τ' -structure where $\tau' = \tau \cup \{Q_i \mid i \in \{0, \dots, p-1\}\}$ and Q_i is a unary relation symbol that does not appear in τ , for each $i \in \{0, \dots, p-1\}$. We have the following lemma.

Lemma 9.3.1. *Let \mathcal{S} be a class of finite τ -structures, and $p, k \in \mathbb{N}$ be such that $p \geq k$. Then the following are true.*

1. \mathcal{L} -EBSP($\mathcal{S}_p, 0$) implies \mathcal{L} -EBSP($\mathcal{S}_k, 0$)
2. \mathcal{L} -EBSP($\mathcal{S}_{p \cdot k}, 0$) implies \mathcal{L} -EBSP($(\mathcal{S}_p)_k, 0$)
3. \mathcal{L} -EBSP($\mathcal{S}_{k+1}, 0$) implies \mathcal{L} -EBSP(\mathcal{S}, k)
4. \mathcal{L} -EBSP(\mathcal{S}, k) implies FO-EBSP(\mathcal{S}, k)

Further, in each of the implications above, any witness function for the antecedent is also a witness function for the consequent.

Proof. Part 1: Obvious since $\mathcal{S}_k \subseteq \mathcal{S}_p$ and \mathcal{S}_k is closed under substructures that are in \mathcal{S}_p .

Part 2: Consider $\mathfrak{A} \in (\mathcal{S}_p)_k$. Each element of \mathfrak{A} can be seen to be labelled with a pair (i, j) of labels where $i \in \{0, \dots, p-1\}$ and $j \in \{0, \dots, k-1\}$. Then \mathfrak{A} can naturally be represented as a structure \mathfrak{A}' of $\mathcal{S}_{p \cdot k}$ as follows: (i) the τ -reducts of \mathfrak{A} and \mathfrak{A}' are the same, and (ii) an element a of \mathfrak{A}' having label (i, j) in \mathfrak{A} is labelled with label $i \times k + j$ in \mathfrak{A}' . Since \mathcal{L} -EBSP($\mathcal{S}_{p \cdot k}, 0$) is true, there exists a witness function $\theta_{(\mathcal{S}_{p \cdot k}, 0, \mathcal{L})} : \mathbb{N} \rightarrow \mathbb{N}$ such that for any $m \in \mathbb{N}$, there exists $\mathfrak{B}' \in \mathcal{S}_{p \cdot k}$ such that (i) $\mathfrak{B}' \subseteq \mathfrak{A}'$ (ii) $|\mathfrak{B}'| \leq \theta_{(\mathcal{S}_{p \cdot k}, 0, \mathcal{L})}(m)$ and (iii) $\mathfrak{B}' \equiv_{m, \mathcal{L}} \mathfrak{A}'$. Consider the structure $\mathfrak{B} \in (\mathcal{S}_p)_k$ such that (i) the τ -reducts of \mathfrak{B}' and \mathfrak{B} are the same, and (ii) if the label of an element a of \mathfrak{B} is $i \times k + j$ in \mathfrak{B}' , then its label in \mathfrak{B} is (i, j) . It is easy to see that (i) $\mathfrak{B} \subseteq \mathfrak{A}$ (ii) $|\mathfrak{B}| \leq \theta_{(\mathcal{S}_{p \cdot k}, 0, \mathcal{L})}(m)$, and (iii) using the same strategy as of the duplicator in the m -round \mathcal{L} -EF game between \mathfrak{B}' and \mathfrak{A}' , the duplicator always wins in the m -round \mathcal{L} -EF game between \mathfrak{B} and \mathfrak{A} ; in other words, $\mathfrak{B} \equiv_{m, \mathcal{L}} \mathfrak{A}$. Then \mathcal{L} -EBSP-condition($(\mathcal{S}_p)_k, \mathfrak{A}, \mathfrak{B}, 0, m, \text{null}, \theta_{((\mathcal{S}_p)_k, 0, \mathcal{L})}$) is true, where null denotes the empty tuple and $\theta_{((\mathcal{S}_p)_k, 0, \mathcal{L})} = \theta_{(\mathcal{S}_{p \cdot k}, 0, \mathcal{L})}$.

Part 3: Let $\mathfrak{A} \in \mathcal{S}$ and $\bar{a} = (a_1, \dots, a_k)$ be a k -tuple from \mathfrak{A} . Consider the structure \mathfrak{A}' of \mathcal{S}_{k+1} whose τ -reduct is \mathfrak{A} and in which the element a_i has been labelled with label $i-1$ for $1 \leq i \leq k$ and all elements other than the a_i s have been labelled with label k . Since \mathcal{L} -EBSP($\mathcal{S}_{k+1}, 0$) holds, there exists a witness function $\theta_{(\mathcal{S}_{k+1}, 0, \mathcal{L})} : \mathbb{N} \rightarrow \mathbb{N}$ such that given any $m \in \mathbb{N}$, there exists $\mathfrak{B}' \in \mathcal{S}_{k+1}$ satisfying (i) $\mathfrak{B}' \subseteq \mathfrak{A}'$ (ii) $|\mathfrak{B}'| \leq \theta_{(\mathcal{S}_{k+1}, 0, \mathcal{L})}(m)$ and (iii) $\mathfrak{B}' \equiv_{m, \mathcal{L}} \mathfrak{A}'$. Let \mathfrak{B} be the τ -reduct of \mathfrak{B}' . It is clear that (i) $\mathfrak{B} \in \mathcal{S}$ since $\mathfrak{B}' \in \mathcal{S}_{k+1}$ (ii) $\mathfrak{B} \subseteq \mathfrak{A}$ (iii) \mathfrak{B} must contain the elements of \bar{a} since a_i is the unique element of \mathfrak{A}' that is labeled with label $i-1$, for each $i \in \{1, \dots, k\}$ (iv) $|\mathfrak{B}| \leq \theta_{(\mathcal{S}_{k+1}, 0, \mathcal{L})}(m)$, and (v) using the same strategy as of the duplicator in the m -round \mathcal{L} -EF game between \mathfrak{B}' and \mathfrak{A}' , the duplicator always wins in the m -round \mathcal{L} -EF game between (\mathfrak{B}, \bar{a}) and (\mathfrak{A}, \bar{a}) ; in other words, $\text{tp}_{\mathfrak{B}, \bar{a}, m, \mathcal{L}}(\bar{x}) = \text{tp}_{\mathfrak{A}, \bar{a}, m, \mathcal{L}}(\bar{x})$. Then \mathcal{L} -EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, k, m, \bar{a}, \theta_{(\mathcal{S}, k, \mathcal{L})}$) is true, where $\theta_{(\mathcal{S}, k, \mathcal{L})} = \theta_{(\mathcal{S}_{k+1}, 0, \mathcal{L})}$.

Part 4: Obvious; since $\text{FO} \subseteq \mathcal{L}$, it follows that for structures (\mathfrak{A}, \bar{a}) and (\mathfrak{B}, \bar{b}) , if $\text{tp}_{\mathfrak{A}, \bar{a}, m, \mathcal{L}}(\bar{x}) = \text{tp}_{\mathfrak{B}, \bar{b}, m, \mathcal{L}}(\bar{x})$, then $\text{tp}_{\mathfrak{A}, \bar{a}, m, \text{FO}}(\bar{x}) = \text{tp}_{\mathfrak{B}, \bar{b}, m, \text{FO}}(\bar{x})$.

We observe that in all the cases above, any witness function for the antecedent is also a witness function for the consequent. \square

In the next chapter, we prove that \mathcal{L} -EBSP(\mathcal{S}, k) holds of several classes \mathcal{S} which are of interest in computer science. In doing so, we use Lemma 9.3.1 in an important way to simplify our proofs since as this lemma shows, to prove \mathcal{L} -EBSP for a class \mathcal{S} and parameter k , it suffices to prove \mathcal{L} -EBSP for the class \mathcal{S}_p and parameter 0, where p is suitably chosen.

Chapter 10

Classes satisfying \mathcal{L} -EBSP(\cdot, k)

In this chapter, we show that various interesting classes of structures satisfy \mathcal{L} -EBSP(\cdot, k), and also give methods to construct new classes of structures that satisfy \mathcal{L} -EBSP(\cdot, k) from classes known to satisfy the latter property. Broadly speaking, the specific classes that we consider are of two kinds – one that are special kinds of posets, and the other that are special kinds of graphs. In Section 10.2, we consider the former kind of classes and prove our results for words, trees (unordered, ordered, or ranked) and nested words over a given finite alphabet Σ . In Section 10.3, we consider the latter kind of classes and prove our results for n -partite cographs, and hence various subclasses of these including cographs, graph classes of bounded tree-depth, graph classes of bounded shrub-depth and graph classes of bounded \mathcal{SC} -depth. In Section 10.4, we show that classes that satisfy \mathcal{L} -EBSP(\cdot, \cdot) are, under suitable assumptions, closed under set-theoretic operations, under operations that are implementable using quantifier-free translation schemes, and under transformations that are defined using regular operation-tree languages, where an operation-tree is a finite composition of the aforementioned operations on classes of structures. These closure properties give us means to construct a wide array of classes that satisfy \mathcal{L} -EBSP(\cdot, \cdot).

All of the above results derive from an abstract result concerning tree representations that we now describe in Section 10.1.

10.1 An abstract result concerning tree representations

An *unlabeled unordered tree* is a finite poset $P = (A, \leq)$ with a unique minimal element (called “root”), and such that for each $c \in A$, the set $\{b \mid b \leq c\}$ is totally ordered by \leq . Informally

speaking, the Hasse diagram of P is an inverted (graph-theoretic) tree. We call A as the set of *nodes* of P . We use the standard notions of leaf, internal node, ancestor, descendent, parent, child, degree, height and subtree in connection with trees. Explicitly, these are as defined below.

Let $P = (A, \leq)$ be a (unlabeled unordered) tree. Let a, b be distinct nodes of A .

1. We say a is a *leaf* of P if for any node $c \in A$, we have $(a \leq c) \rightarrow (c = a)$. A node of P that is not a leaf of P is called an *internal node* of P .
2. We say a is an *ancestor* of b in P , or equivalently that b is a *descendent* of a in P , if $a \leq b$ and $a \neq b$. A *common ancestor* of a and b in P is a node c of P such that $c \leq a$ and $c \leq b$. The *greatest common ancestor* of a and b in P , denoted $a \wedge_P b$, is a common ancestor of a and b in P such that for every common ancestor c of a and b in P , we have $c \leq (a \wedge_P b)$.
3. We say a is the *parent* of b in P , or equivalently, that b is a *child* of a in P , if a is an ancestor of b in P and any ancestor of b in P is either a itself or an ancestor of a in P . We let $\text{Children}_P(a)$ denote the set $\{b \mid b \text{ is a child of } a \text{ in } P\}$.
4. The *degree* of a in P is the size of $\text{Children}_P(a)$. The degree of P is the maximum of the degrees of the nodes of P .
5. The *height* of P is one less than the size of the longest chain in P .
6. The *subtree of P induced by a subset A' of A* is the tree (A', \leq') where $\leq' = \leq \cap (A' \times A')$. A *subtree of P* is a subtree of P induced by some subset of A .

An *unlabeled ordered tree* is a pair $O = (P, \lesssim)$ where P is an unlabeled unordered tree and \lesssim is a binary relation that imposes a linear order on the children of any internal node of P . In this section, by trees we always mean *ordered trees*. It is clear that the notions introduced above for unordered trees can be adapted for ordered trees. We define some additional notions for ordered trees below. It is clear that these notions can be adapted for unordered trees.

7. Given a countable alphabet Σ , a *tree over Σ* , also called a Σ -*tree*, or simply *tree* when Σ is clear from context, is a pair (O, λ) where O is an unlabeled tree and $\lambda : A \rightarrow \Sigma$ is a labeling function, where A is the set of nodes of O . We denote Σ -trees by s, t, x, y or z , possibly with numbers as subscripts.
8. Given a tree t , we denote the root of t as $\text{root}(t)$. For a node a of t , we denote the subtree of t rooted at a as $t_{\geq a}$, and the subtree of t obtained by deleting $t_{\geq a}$ from t , as $t - t_{\geq a}$.
9. Given a tree s and a non-root node a of t , the *replacement of $t_{\geq a}$ with s in t* , denoted $t[t_{\geq a} \mapsto s]$, is a tree defined as follows: let c be the parent of a in t and s' be an isomorphic copy of s whose nodes are disjoint with those of t . Then $t[t_{\geq a} \mapsto s]$ is defined upto

isomorphism as the tree obtained by deleting $t_{\geq a}$ from t to get a tree t' , and inserting (the root of) s' at the same position among the children of c in t' , as the position of a among the children of c in t .

10. For t, s and s' as in the previous point, suppose the roots of each of these trees have the same label. Then the *merge of s with t* , denoted $t \odot s$, is defined upto isomorphism as the tree obtained by deleting $\text{root}(s')$ from s' and concatenating the sequence of subtrees hanging at $\text{root}(s')$ in s' , to the sequence of subtrees hanging at $\text{root}(t)$ in t . Thus the children of $\text{root}(s')$ in s' are the “new” children of $\text{root}(t)$, and appear “after” the “old” children of $\text{root}(t)$, and in the order they appear in s' .

Fix a finite alphabet Σ_{int} and a countable alphabet Σ_{leaf} (where the two alphabets are allowed to be overlapping). We say a class \mathcal{T} of $(\Sigma_{\text{int}} \cup \Sigma_{\text{leaf}})$ -trees is *representation-feasible* if it is closed under (label-preserving) isomorphisms, and if every tree $t = (O, \lambda)$ in the class has the property that for every leaf, resp. internal, node a of t , the label $\lambda(a)$ belongs to Σ_{leaf} , resp. Σ_{int} . Given a class \mathcal{S} of τ -structures, let $\text{Str} : \mathcal{T} \rightarrow \mathcal{S}$ be a map that associates with each tree in \mathcal{T} , a structure in \mathcal{S} . For a tree $t \in \mathcal{T}$, if $\mathfrak{A} = \text{Str}(t)$, then we say t is a *tree representation* of \mathfrak{A} under Str , or simply a *tree representation* of \mathfrak{A} . We call Str as a *representation map*, and \mathcal{T} as a class of representation trees. For the purposes of our result, we consider maps Str that map isomorphic (preserving labels) trees to isomorphic structures, and that have additional properties among those mentioned below.

A. Transfer properties:

1. Let $t, s_1 \in \mathcal{T}$, t be of size ≥ 2 , and a be a child of $\text{root}(t)$. Suppose $s_2 = t_{\geq a}$ and $z = t[s_2 \mapsto s_1] \in \mathcal{T}$.
 - a. If $\text{Str}(s_1) \hookrightarrow \text{Str}(s_2)$, then $\text{Str}(z) \hookrightarrow \text{Str}(t)$.
 - b. If $\text{Str}(s_1) \equiv_{m, \mathcal{L}} \text{Str}(s_2)$, then $\text{Str}(z) \equiv_{m, \mathcal{L}} \text{Str}(t)$.
2. Let $t, s_1, s_2 \in \mathcal{T}$ be trees of size ≥ 2 such that the roots of all these trees have the same label. For $i \in \{1, 2\}$, suppose $z_i = t \odot s_i \in \mathcal{T}$. If $\text{Str}(s_1) \equiv_{m, \mathcal{L}} \text{Str}(s_2)$, then $\text{Str}(z_1) \equiv_{m, \mathcal{L}} \text{Str}(z_2)$.

B. Monotonicity: Let $t \in \mathcal{T}$ and a be a child of $\text{root}(t)$.

1. If $s = t_{\geq a} \in \mathcal{T}$, then $\text{Str}(s) \hookrightarrow \text{Str}(t)$
2. If $s = (t - t_{\geq a}) \in \mathcal{T}$, then $\text{Str}(s) \hookrightarrow \text{Str}(t)$.

We say a representation map is *\mathcal{L} -height-reduction favourable* if it satisfies conditions A.1.a, A.1.b and B.1 above, for all $m \geq m_0$, for some $m_0 \in \mathbb{N}$. A representation map is said to be

\mathcal{L} -degree-reduction favourable if it satisfies all the conditions above, except possibly B.1, for all $m \geq m_0$, for some $m_0 \in \mathbb{N}$. The following result justifies why the two kinds of representation maps just defined, are called so. This result contains the core argument of most results in the subsequent sections of this chapter. Below, \mathcal{T} is said to be *closed under rooted subtrees and under replacements with rooted subtrees* if for all $t \in \mathcal{T}$ and non-root nodes a and b of t such that a is an ancestor of b in t , we have that each of the subtrees $t_{\geq a}$ and $t[t_{\geq a} \mapsto t_{\geq b}]$ are in \mathcal{T} .

Theorem 10.1.1. *For $I = \{1, \dots, n\}$ and $i \in I$, let \mathcal{S}_i be a class of τ_i -structures, $\text{Str}_i : \mathcal{T} \rightarrow \mathcal{S}_i$ be a representation map, and \mathcal{L}_i be either FO or MSO. Then there exist computable functions $\eta_1, \eta_2 : \mathbb{N}^n \rightarrow \mathbb{N}$, such that for each $t \in \mathcal{T}$ and $m_1, \dots, m_n \in \mathbb{N}$, we have the following.*

1. *If \mathcal{T} is closed under subtrees, and Str_i is \mathcal{L}_i -degree-reduction favourable for all $i \in I$, then there exists a subtree s_1 of t in \mathcal{T} , of degree at most $\eta_1(m_1, \dots, m_n)$, such that for all $i \in I$ (i) $\text{Str}_i(s_1) \hookrightarrow \text{Str}_i(t)$, and (ii) $\text{Str}_i(s_1) \equiv_{m_i, \mathcal{L}_i} \text{Str}_i(t)$.*
2. *If \mathcal{T} is closed under rooted subtrees and under replacements with rooted subtrees, and Str_i is \mathcal{L}_i -height-reduction favourable for all $i \in I$, then there exists a subtree s_2 of t in \mathcal{T} , of height at most $\eta_2(m_1, \dots, m_n)$, such that for all $i \in I$ (i) $\text{Str}_i(s_2) \hookrightarrow \text{Str}_i(t)$, and (ii) $\text{Str}_i(s_2) \equiv_{m_i, \mathcal{L}_i} \text{Str}_i(t)$.*

Proof. We recall from Section 7.2 of Chapter 7 that for a class \mathcal{S} of structures, $\Delta_{\mathcal{L}}(m, \mathcal{S})$ denotes the set of all equivalence classes of the $\equiv_{m, \mathcal{L}}$ relation restricted to the structures in \mathcal{S} , and $\Lambda_{\mathcal{S}, \mathcal{L}} : \mathbb{N} \rightarrow \mathbb{N}$ is a fixed computable function with the property that $\Lambda_{\mathcal{S}, \mathcal{L}}(m) \geq |\Delta_{\mathcal{L}}(m, \mathcal{S})|$.

(Part 1): For $i \in I$, let k_i be such that Str_i satisfies all the transfer and monotonicity properties, except possibly B.1, for all $m \geq k_i$. Define $\eta_1 : \mathbb{N}^n \rightarrow \mathbb{N}$ as follows: for $l_1, \dots, l_n \in \mathbb{N}$, $\eta_1(l_1, \dots, l_n) = \prod_{i \in I} \Lambda_{\mathcal{S}_i, \mathcal{L}_i}(\max(k_i, l_i))$. Then η_1 is computable. Now given $m_1, \dots, m_n \in \mathbb{N}$, let $p = \eta_1(m_1, \dots, m_n)$. If t has degree $\leq p$, then putting $s_1 = t$ we are done. Else, some node of t , say a , has degree $r > p$. Let $z = t_{\geq a}$. Let a_1, \dots, a_r be the sequence of children of root(z) in z . For $j \in \{1, \dots, r\}$, let x_j , resp. y_j , be the subtree of z obtained from z by deleting the subtrees rooted at a_j, a_{j+1}, \dots, a_r , resp. deleting the subtrees rooted at a_1, a_2, \dots, a_{j-1} . Then $z = y_1 = x_j \odot y_j$ for all $j \in \{2, \dots, r\}$. Let $q_i = \max(k_i, m_i)$ for $i \in I$ and let $g : \{1, \dots, r\} \rightarrow \prod_{i \in I} \Delta_{\mathcal{L}_i}(q_i, \mathcal{S}_i)$ be such that for $j \in \{1, \dots, r\}$, $g(j)$ is the sequence $(\delta_i)_{i \in I}$ where δ_i is the $\equiv_{q_i, \mathcal{L}_i}$ class of $\text{Str}_i(y_j)$. Verify that $|\prod_{i \in I} \Delta_{\mathcal{L}_i}(q_i, \mathcal{S}_i)| \leq p$. Then since $r > p$, there exist $j, k \in \{1, \dots, r\}$ such that $j < k$ and $g(j) = g(k)$, i.e. for all $i \in I$, $\text{Str}_i(y_j) \equiv_{q_i, \mathcal{L}_i} \text{Str}_i(y_k)$. If $z_1 = x_j \odot y_k$, then since \mathcal{T} is closed under subtrees, we have $z_1 \in \mathcal{T}$. By the properties B.2 and A.2 above, we have $\text{Str}_i(z_1) \hookrightarrow \text{Str}_i(z)$ and $\text{Str}_i(z_1) \equiv_{q_i, \mathcal{L}_i} \text{Str}_i(z)$, for all $i \in I$. By iteratively

applying properties A.1.a and A.1.b to the nodes along the path from a to $\text{root}(t)$, we see that if $t_1 = t[z \mapsto z_1]$, then $\text{Str}_i(t_1) \hookrightarrow \text{Str}_i(t)$ and $\text{Str}_i(t_1) \equiv_{q_i, \mathcal{L}_i} \text{Str}_i(t)$, for all $i \in I$. Observe that t_1 has strictly lesser size than t . Recursing on t_1 , we eventually get a subtree s_1 of t of degree at most p such that for all $i \in I$, (i) $\text{Str}_i(s_1) \hookrightarrow \text{Str}_i(t)$ and (ii) $\text{Str}_i(s_1) \equiv_{q_i, \mathcal{L}_i} \text{Str}_i(t)$. Since $q_i = \max(k_i, m_i)$, we have $\text{Str}_i(s_1) \equiv_{m_i, \mathcal{L}_i} \text{Str}_i(t)$ for all $i \in I$.

(Part 2): For $i \in I$, let r_i be such that Str_i satisfies conditions A.1.a, A.1.b and B.1, for all $m \geq r_i$. Define $\eta_2 : \mathbb{N}^n \rightarrow \mathbb{N}$ as follows: for $l_1, \dots, l_n \in \mathbb{N}$, $\eta_1(l_1, \dots, l_n) = \prod_{i \in I} \Delta_{\mathcal{L}_i}(q_i, \mathcal{S}_i)$. Then η_2 is computable. Now given $m_1, \dots, m_n \in \mathbb{N}$, let $p = \eta_2(m_1, \dots, m_n)$. If t has height $\leq p$, then putting $s_2 = t$ we are done. Else there is a path from the root of t to some leaf of t , whose length is $> p$. Let A be the set of nodes appearing along this path. For $i \in I$, let $q_i = \max(r_i, m_i)$. Consider the function $h : A \rightarrow \prod_{i \in I} \Delta_{\mathcal{L}_i}(q_i, \mathcal{S}_i)$ such that for each $a \in A$, $h(a) = (\delta_i)_{i \in I}$ where δ_i is the $\equiv_{q_i, \mathcal{L}_i}$ class of $\text{Str}_i(t_{\geq a})$. Verify that $|\prod_{i \in I} \Delta_{\mathcal{L}_i}(q_i, \mathcal{S}_i)| \leq p$. Since $|A| > p$, there exist distinct nodes $a, b \in A$ such that a is an ancestor of b in t and $h(a) = h(b)$. We have two cases: (i) Node a is the root of t ; then let $t_2 = t_{\geq b}$. (ii) Node a is not the root of t ; then let $t_2 = t[t_{\geq a} \mapsto t_{\geq b}]$. Since \mathcal{T} is closed under rooted subtrees and under replacements with rooted subtrees, we have $t_{\geq a}, t_{\geq b}$ and t_2 are all in \mathcal{T} . By property B.1, $\text{Str}_i(t_{\geq b}) \hookrightarrow \text{Str}_i(t_{\geq a})$. Also since $h(a) = h(b)$, we have $\text{Str}_i(t_{\geq b}) \equiv_{q_i, \mathcal{L}_i} \text{Str}_i(t_{\geq a})$. Then by iteratively applying properties A.1.a and A.1.b to the nodes along the path from a to $\text{root}(t)$, we get in either of the cases above, that $\text{Str}_i(t_2) \hookrightarrow \text{Str}_i(t)$ and $\text{Str}_i(t_2) \equiv_{q_i, \mathcal{L}_i} \text{Str}_i(t)$, for all $i \in I$. Observe that t_2 has strictly less size than t . Recursing on t_2 , we eventually get a subtree s_2 of t of height at most p such that for all $i \in I$, (i) $\text{Str}_i(s_2) \hookrightarrow \text{Str}_i(t)$ and (ii) $\text{Str}_i(s_2) \equiv_{q_i, \mathcal{L}_i} \text{Str}_i(t)$. Since $q_i = \max(r_i, m_i)$, we have $\text{Str}_i(s_2) \equiv_{m_i, \mathcal{L}_i} \text{Str}_i(t)$ for all $i \in I$. \square

Call a representation map $\text{Str} : \mathcal{T} \rightarrow \mathcal{S}_1$ as *size effective* if there is a computable function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that $|\text{Str}(t)| \leq f(|t|)$ for all $t \in \mathcal{T}$. Call Str *onto upto isomorphism* if for every structure in \mathcal{S}_1 , there is an isomorphic structure that is in the range of Str . For a given class \mathcal{S} , we say \mathcal{S} admits an \mathcal{L} -reduction favourable size effective (\mathcal{L} -RFSE) tree representation if there exist finite alphabets Σ_{int} and Σ_{leaf} , a class \mathcal{T} of representation-feasible trees over $\Sigma_{\text{int}} \cup \Sigma_{\text{leaf}}$, and a size effective representation map $\text{Str} : \mathcal{T} \rightarrow \mathcal{S}$ that is onto upto isomorphism, such that \mathcal{T} and Str are of one of the following types:

- \mathcal{L} -RFSE-type I: \mathcal{T} has bounded degree and is closed under rooted subtrees and under replacements with rooted subtrees, and Str is \mathcal{L} -height-reduction favourable.
- \mathcal{L} -RFSE-type II: \mathcal{T} is closed under subtrees, and Str is both \mathcal{L} -height-reduction favourable

and \mathcal{L} -degree-reduction favourable.

We say \mathcal{S} admits an \mathcal{L} -RFSE tree representation schema if for each $p \in \mathbb{N}$, the class \mathcal{S}_p (defined in Section 9.3) admits an \mathcal{L} -RFSE tree representation. We now have the following result.

Lemma 10.1.2. *Let \mathcal{S} be a class of structures that admits an \mathcal{L} -RFSE tree representation schema. Then \mathcal{L} -EBSP(\mathcal{S}_p, k) holds with a computable witness function, for all $p, k \in \mathbb{N}$,*

Proof. By Lemma 9.3.1, it suffices to show that \mathcal{L} -EBSP($\mathcal{S}_r, 0$) holds for $r = p \cdot (k + 1)$, with a computable witness function. Since \mathcal{S} admits an \mathcal{L} -RFSE tree representation schema, the class \mathcal{S}_r admits an \mathcal{L} -RFSE tree representation. Let the latter fact be witnessed by a representation-feasible class \mathcal{T} of trees, and a representation map $\text{Str} : \mathcal{T} \rightarrow \mathcal{S}_r$. We have two cases from the definition of \mathcal{L} -RFSE tree representation:

1. \mathcal{T} and Str are of \mathcal{L} -RFSE-type I: Given $\mathfrak{A} \in \mathcal{S}_r$, let $t \in \mathcal{T}$ be such that $\text{Str}(t) \cong \mathfrak{A}$ (this is guaranteed since Str is onto upto isomorphism). Let $m \in \mathbb{N}$. By Theorem 10.1.1, there is a computable function $\eta_2 : \mathbb{N} \rightarrow \mathbb{N}$ and a subtree s of t in \mathcal{T} such that (i) the height of s is at most $h = \eta_2(m)$, (ii) $\text{Str}(s) \hookrightarrow \text{Str}(t)$ and (iii) $\text{Str}(s) \equiv_{m, \mathcal{L}} \text{Str}(t)$. Since the degree of s is bounded, by say d , the size of s is at most d^{h+1} . Whereby if f is the computable function witnessing the size effectiveness of Str , then $|\text{Str}(s)| \leq f(d^h)$. Let \mathfrak{B} be the substructure of \mathfrak{A} such that $\mathfrak{B} \cong \text{Str}(s)$. Since \mathcal{S} is closed under isomorphisms, so is \mathcal{S}_r , whereby $\mathfrak{B} \in \mathcal{S}_r$. We can now see that \mathcal{L} -EBSP-condition($\mathcal{S}_r, \mathfrak{A}, \mathfrak{B}, 0, m, \text{null}, \theta_{(\mathcal{S}_r, 0, \mathcal{L})}$) is true, where null is the empty tuple and $\theta_{(\mathcal{S}_r, 0, \mathcal{L})}(m) = f(d^h)$. Clearly $\theta_{(\mathcal{S}_r, 0, \mathcal{L})}$ is computable.
2. \mathcal{T} and Str are of \mathcal{L} -RFSE-type II: By using both parts of Theorem 10.1.1 and reasoning similarly as in the previous case, we can show that \mathcal{L} -EBSP($\mathcal{S}_r, 0$) holds with a computable witness function given by $\theta_{(\mathcal{S}_r, 0, \mathcal{L})}(m) = f(d^{h+1})$ where $d = \eta_1(m)$ and $h = \eta_2(m)$. Here, f is the computable function witnessing the size-effectiveness of Str , and η_1, η_2 are the computable functions given by Theorem 10.1.1.

□

10.2 Words, trees and nested words

Let Σ be a finite alphabet. The notion of unordered and ordered Σ -trees was already introduced in the previous section. A Σ -tree whose underlying poset is a linear order is called a Σ -word. An ordered Σ -tree $t = (((A, \leq), \lesssim), \lambda)$ is said to be *ranked* by a function $\rho : \Sigma \rightarrow \mathbb{N}$ if the number of children of any internal node a of t is exactly $\rho(\lambda(a))$.

Nested words were introduced by Alur and Madhusudan in [5]. Intuitively speaking, a nested word is a Σ -word equipped with a binary relation that is interpreted as a matching. Formally, given a finite alphabet Σ , a *nested word over Σ* , henceforth also called a *nested Σ -word*, is a 4-tuple $(A, \leq, \lambda, \rightsquigarrow)$, where A is a finite set (called the *set of positions*), \leq is a total linear order on A , $\lambda : A \rightarrow \Sigma$ is a labeling function, and \rightsquigarrow is a binary *matching relation* on A . Each pair $(i, j) \in \rightsquigarrow$ is called a *nesting edge*, the position i is called a *call position*, and the position j is called a *return successor*. The relation \rightsquigarrow satisfies the following properties. Below, $i < j$ denotes $((i \leq j) \wedge (i \neq j))$.

1. Nesting edges go only forward: For $i, j \in A$, if $i \rightsquigarrow j$, then $i < j$.
2. No two nesting edges share a position: For $i \in A$, each of the sets $\{j \in A \mid i \rightsquigarrow j\}$ and $\{j \in A \mid j \rightsquigarrow i\}$ has cardinality at most 1.
3. Nesting edges do not cross: For $i_1, i_2, j_1, j_2 \in A$, if $i_1 \rightsquigarrow j_1$ and $i_2 \rightsquigarrow j_2$, then it is not the case that $i_1 < i_2 \leq j_1 < j_2$.

Remark 10.2.1. The definition of nested words presented above corresponds to the definition of nested words in [5], in which the nested words *do not have any pending calls or pending returns*. Hence, the elements $+\infty$ and $-\infty$ present in the definition in [5] are not necessary here, and have been dropped in the definition above without any loss of generality.

For each of the classes of words, trees (unordered, ordered and ranked) and nested words introduced above, the notion of *regularity* of a subclass is well-studied in the literature. For words and nested words, this notion is defined in terms of finite state (word) automata and finite state nested word automata [5]. For each of the aforementioned classes of trees, regularity is defined in terms of variants of finite state tree automata [13]. All of these notions of regularity have been shown to be equivalent to definability via MSO sentences [5, 13]. Therefore, in our result below, a subclass \mathcal{S} of any of the classes \mathcal{S}' of words, trees and nested words, is said to be *regular* if it is definable over \mathcal{S}' using an MSO sentence (in other words, \mathcal{S} is the class of models in \mathcal{S}' , of an MSO sentence). The central theorem of this section is now stated as follows.

Theorem 10.2.2. *Given a finite alphabet Σ and a function $\rho : \Sigma \rightarrow \mathbb{N}$, let $\text{Words}(\Sigma)$, $\text{Unordered-trees}(\Sigma)$, $\text{Ordered-trees}(\Sigma)$, $\text{Ordered-ranked-trees}(\Sigma, \rho)$ and $\text{Nested-words}(\Sigma)$ denote respectively, the classes of all Σ -words, all unordered Σ -trees, all ordered Σ -trees, all ordered Σ -trees ranked by ρ , and all nested Σ -words. Let \mathcal{S} be a regular subclass of any of these classes. Then $\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ holds with a computable witness function for each $k \in \mathbb{N}$.*

We devote the rest of this section to proving Theorem 10.2.2.

Proof of Theorem 10.2.2 for words and trees

We show $\text{MSO-EBSP}(\mathcal{S}, k)$ holds with a computable witness function when \mathcal{S} is exactly one of the classes $\text{Words}(\Sigma)$, $\text{Unordered-trees}(\Sigma)$, $\text{Ordered-trees}(\Sigma)$, and $\text{Ordered-ranked-trees}(\Sigma, \rho)$. That \mathcal{L} -EBSP(\cdot, k) holds with a computable witness function for a regular subclass follows, because (i) a regular subclass of any of the above classes is MSO definable over the class, (ii) $\text{MSO-EBSP}(\cdot, k)$ and the computability of witness function are preserved under MSO definable subclasses (Lemma 10.4.1(4)), and (iii) $\text{MSO-EBSP}(\cdot, k)$ implies $\text{FO-EBSP}(\cdot, k)$, and any witness function for the former is also a witness function for the latter (Lemma 9.3.1(4)).

Of the classes mentioned above, we show $\text{MSO-EBSP}(\mathcal{S}, k)$ holds with a computable witness function for the case when \mathcal{S} is either $\text{Unordered-trees}(\Sigma)$ or $\text{Ordered-ranked-trees}(\Sigma, \rho)$. The proof for $\text{Ordered-trees}(\Sigma)$ can be done similarly. That the result holds for $\text{Words}(\Sigma)$ follows from the fact that $\text{Words}(\Sigma)$ is a subclass of $\text{Unordered-trees}(\Sigma)$ that is hereditary over the latter, and then by using Lemma 10.4.1(1).

For our proofs, we need MSO *composition lemmas* for unordered trees and ordered trees. Composition results were first studied by Feferman and Vaught, and subsequently by many others (see [57]). We state the MSO composition lemma first for ordered trees, towards which we define some terminology. For a finite alphabet Ω , given ordered Ω -trees t, s and a non-root node a of t , the *join of s to t to the right of a* , denoted $t \cdot_a^{\rightarrow} s$, is defined as follows: Let s' be an isomorphic copy of s whose set of nodes is disjoint with the set of nodes of t . Then $t \cdot_a^{\rightarrow} s$ is defined upto isomorphism as the tree obtained by making s' as a new child subtree of the parent of a in t , at the successor position of the position of a among the children of t . We can similarly define the *join of s to t to the left of a* , denoted $t \cdot_a^{\leftarrow} s$. Likewise, for t and s as above, if a is a leaf node of t , we can define the *join of s to t below a* , denoted $t \cdot_a^{\uparrow} s$, as the tree obtained upto isomorphism by making the root of s as a child of a . The MSO composition lemma for ordered trees can now be stated as follows. The proof of this lemma is provided towards the end of this section.

Lemma 10.2.3 (Composition lemma for ordered trees). *For a finite alphabet Ω , let t_i, s_i be non-empty ordered Ω -trees, and let a_i be a non-root node of t_i , for each $i \in \{1, 2\}$. Let $m \geq 2$ and suppose that $(t_1, a_1) \equiv_{m, \text{MSO}} (t_2, a_2)$ and $s_1 \equiv_{m, \text{MSO}} s_2$. Then each of the following hold.*

- I. $((t_1 \cdot_{a_1}^{\rightarrow} s_1), a_1) \equiv_{m, \text{MSO}} ((t_2 \cdot_{a_2}^{\rightarrow} s_2), a_2)$

2. $((t_1 \cdot_{a_1}^{\leftarrow} s_1), a_1) \equiv_{m,MSO} ((t_2 \cdot_{a_2}^{\leftarrow} s_2), a_2)$
3. $((t_1 \cdot_{a_1}^{\uparrow} s_1), a_1) \equiv_{m,MSO} ((t_2 \cdot_{a_2}^{\uparrow} s_2), a_2)$ if a_1, a_2 are leaf nodes of t_1, t_2 resp.

We now state the MSO composition lemma for unordered trees, towards which we introduce terminology akin to that introduced above for ordered trees. Given unordered trees t and s , and a node a of t , define the *join of s to t to a* , denoted $t \cdot_a s$, as follows: Let s' be an isomorphic copy of s whose set of nodes is disjoint with the set of nodes of t . Then $t \cdot_a s$ is defined upto isomorphism as the tree obtained by making s' as a new child subtree of a in t . The MSO composition lemma for unordered trees is now as stated below. The proof is similar to that of Lemma 10.2.3, and is hence skipped.

Lemma 10.2.4 (Composition lemma for unordered trees). *For a finite alphabet Ω , let t_i, s_i be non-empty unordered Ω -trees, and let a_i be a node of t_i , for each $i \in \{1, 2\}$. For $m \in \mathbb{N}$, suppose that $(t_1, a_1) \equiv_{m,MSO} (t_2, a_2)$ and $s_1 \equiv_{m,MSO} s_2$. Then $((t_1 \cdot_{a_1} s_1), a_1) \equiv_{m,MSO} ((t_2 \cdot_{a_2} s_2), a_2)$.*

We now show \mathcal{L} -EBSP(\cdot, k) holds with a computable witness function for each of the classes Unordered-trees(Σ) and Ordered-ranked-trees(Σ, ρ).

1. Let \mathcal{S} be the class of all unlabeled unordered trees. We show that \mathcal{S} admits an MSO-RFSE tree representation schema. Then using Lemma 10.1.2, we get that MSO-EBSP(\mathcal{S}_p, k) holds with a computable witness function for each $p \in \mathbb{N}$. Since there is a 1-1 correspondence between \mathcal{S}_p and Unordered-trees(Σ) when $|\Sigma| = p$, it follows that MSO-EBSP(Unordered-trees(Σ), k) holds with a computable witness function.

Consider the class \mathcal{S}_p where $p \geq 1$. Let \mathcal{T} be the class of all representation-feasible trees over $\Sigma_{\text{int}} \cup \Sigma_{\text{leaf}}$ where $\Sigma_{\text{int}} = \Sigma_{\text{leaf}} = \{0, \dots, p-1\}$. There is a natural map $\text{Str} : \mathcal{T} \rightarrow \mathcal{S}_p$ that simply “forgets” the ordering among the children of any node of its input tree. More precisely, for an ordered tree (O, λ) over $\{0, \dots, p-1\}$ where $O = ((A, \leq), \lesssim)$, we have $\text{Str}((O, \lambda)) = ((A, \leq), \lambda)$. It is easy to see that Str satisfies for $\mathcal{L} = \text{MSO}$, the conditions A.1.a, B.1 and B.2 stated in Section 10.1. That Str satisfies for $\mathcal{L} = \text{MSO}$, the conditions A.1.b and A.2 for $m \geq 0$ follows from Lemma 10.2.4 above. Then Str is MSO-height-reduction favourable and MSO-degree-reduction favourable. That \mathcal{T} is closed under subtrees, and that Str is size effective and onto upto isomorphism, are obvious. Then \mathcal{T} and Str are of MSO-RFSE-type II. Since p is arbitrary, we get that \mathcal{S} admits an MSO-RFSE tree representation schema.

2. Let \mathcal{S} be the class of all unlabeled ordered trees ranked by ρ . We show that \mathcal{S} admits an MSO-RFSE tree representation schema. Then by Lemma 10.1.2, we get that $\text{MSO-EBSP}(\mathcal{S}_p, k)$ holds with a computable witness function for each $p \in \mathbb{N}$. Since there is a 1-1 correspondence between \mathcal{S}_p and $\text{Ordered-ranked-trees}(\Sigma, \rho)$ when $|\Sigma| = p$, it follows that $\text{MSO-EBSP}(\text{Ordered-ranked-trees}(\Sigma, \rho), k)$ holds with a computable witness function. Consider the class \mathcal{S}_p where $p \geq 1$. Let $\Sigma_{\text{int}} = \Sigma_{\text{leaf}} = \{0, \dots, p-1\}$, and let \mathcal{T} be the class of all representation-feasible trees over $\Sigma_{\text{int}} \cup \Sigma_{\text{leaf}}$, that are ranked by ρ . Indeed, then $\mathcal{T} = \mathcal{S}_p$. That \mathcal{T} is of bounded degree, and is closed under rooted subtrees and under replacements with rooted subtrees is clear. Let $\text{Str} : \mathcal{T} \rightarrow \mathcal{S}_p$ be the identity map. That Str satisfies for $\mathcal{L} = \text{MSO}$, the conditions A.1.a and B.1, and that Str is size effective and onto upto isomorphism, are clear. That Str satisfies A.1.b for $m \geq 2$ follows from Lemma 10.2.3. Whence Str is MSO-height-reduction favourable, whereby \mathcal{T} and Str are of MSO-RFSE-type I. Since p is arbitrary, we get that \mathcal{S} admits an MSO-RFSE tree representation schema.

We now prove the MSO composition lemma for ordered trees.

Proof of Lemma 10.2.3. Without loss of generality, we assume \mathbf{t}_i and \mathbf{s}_i have disjoint sets of nodes for $i \in \{1, 2\}$. We show the result for part (1) above. The others are similar. Let $\mathbf{z}_i = (\mathbf{t}_i \cdot_{a_i} \mathbf{s}_i)$ for $i \in \{1, 2\}$.

Let β_1 be the winning strategy of the duplicator in the m round MSO-EF game between (\mathbf{t}_1, a_1) and (\mathbf{t}_2, a_2) . Let β_2 be the winning strategy of the duplicator in the m round MSO-EF game between \mathbf{s}_1 and \mathbf{s}_2 . Observe that since $m \geq 2$, we have β_2 is such that if the spoiler picks $\text{root}(\mathbf{s}_1)$ (resp. $\text{root}(\mathbf{s}_2)$), then β_2 will require the duplicator to pick $\text{root}(\mathbf{s}_2)$ (resp. $\text{root}(\mathbf{s}_1)$). We use this observation later on. The strategy α of the duplicator in the m -round MSO-EF game between (\mathbf{z}_1, a_1) and (\mathbf{z}_2, a_2) is defined as follows:

1. Point move: (i) If the spoiler picks an element of \mathbf{t}_1 (resp. \mathbf{t}_2), the duplicator picks the element of \mathbf{t}_2 (resp. \mathbf{t}_1) given by β_1 . (ii) If the spoiler picks an element of \mathbf{s}_1 (resp. \mathbf{s}_2), the duplicator picks the element of \mathbf{s}_2 (resp. \mathbf{s}_1) given by β_2 .
2. Set move: If the spoiler picks a set X from \mathbf{z}_1 , then let $X = Y_1 \sqcup Y_2$ where Y_1 is a set of elements of \mathbf{t}_1 and Y_2 is a set of elements of \mathbf{s}_1 . Let Y'_1 and Y'_2 be the sets of elements of \mathbf{t}_2 and \mathbf{s}_2 respectively, chosen according to strategies β_1 and β_2 . Then in the game between (\mathbf{z}_1, a_1) and (\mathbf{z}_2, a_2) , the duplicator responds with the set $X' = Y'_1 \cup Y'_2$. A similar choice of set is made by the duplicator from \mathbf{z}_1 when the spoiler chooses a set from \mathbf{z}_2 .

We now show that the strategy α is winning for the duplicator in the m -round MSO-EF game between (z_1, a_1) and (z_2, a_2) .

Let at the end of m rounds, the vertices and sets chosen from z_1 , resp. z_2 , be e_1, \dots, e_p and E_1, \dots, E_r , resp. f_1, \dots, f_p and F_1, \dots, F_r , where $p + r = m$. For $l \in \{1, \dots, r\}$, let E_l^t , resp. E_l^s be the intersection of E_l with the nodes of t_1 , resp. nodes of s_1 , and likewise, let F_l^t , resp. F_l^s be the intersection of F_l with the nodes of t_2 , resp. nodes of s_2 .

Firstly, it is straightforward to verify that the labels of e_i and f_i are the same for all $i \in \{1, \dots, p\}$, and that for $l \in \{1, \dots, r\}$, e_i is in E_l^s , resp. E_l^t , iff f_i is in F_l^s , resp. F_l^t , whereby $e_i \in E_l$ iff $f_i \in F_l$. For $1 \leq i, j \leq p$, if e_i and e_j both belong to t_1 or both belong to s_1 , then it is clear from the strategy α described above, that f_i and f_j both belong resp. to t_2 or both belong to s_2 . It is easy to verify from the description of α that for every binary relation (namely, the ancestor-descendent-order \leq , and the ordering-on-the-children-order \lesssim), the pair (e_i, e_j) is in the binary relation in z_1 iff (f_i, f_j) is in that binary relation in z_2 . Consider the case when without loss of generality, $e_1 \in t_1$ and $e_2 \in s_1$. Then $f_1 \in t_2$ and $f_2 \in s_2$. We have the following cases. Assume that the ordered tree underlying z_i is $((A_i, \leq_i), \lesssim_i)$ for $i \in \{1, 2\}$.

1. $e_1 \lesssim_1 a_1$ and $e_2 = \text{root}(s_1)$: Then we see that $f_1 \lesssim_2 a_2$ and $f_2 = \text{root}(s_2)$. Observe that f_2 must be $\text{root}(s_2)$ by the property of β_2 stated at the outset. Whereby $e_1 \lesssim_1 e_2$ and $f_1 \lesssim_2 f_2$. Likewise $e_1 \not\lesssim_1 e_2$, $e_2 \not\lesssim_1 e_1$ and $f_1 \not\lesssim_2 f_2$, $f_2 \not\lesssim_2 f_1$.
2. $e_1 \lesssim_1 a_1$ and $e_2 \neq \text{root}(s_1)$: Then we see that $f_1 \lesssim_2 a_2$ and $f_2 \neq \text{root}(s_2)$ (again by the property of β_2 stated at the outset). Whereby $e_1 \lesssim_1 e_2$, $e_2 \lesssim_1 e_1$ and $f_1 \lesssim_2 f_2$, $f_2 \lesssim_2 f_1$. Likewise, $e_1 \not\lesssim_1 e_2$, $e_2 \not\lesssim_1 e_1$ and $f_1 \not\lesssim_2 f_2$, $f_2 \not\lesssim_2 f_1$.
3. $a_1 \lesssim_1 e_1$, $a_1 \neq e_1$ and $e_2 = \text{root}(s_1)$: Then we see that $a_2 \lesssim_2 f_1$, $a_2 \neq f_1$ and $f_2 = \text{root}(s_2)$. Observe that f_2 must be $\text{root}(s_2)$ by the property of β_2 stated at the outset. Whereby $e_2 \lesssim_1 e_1$ and $f_2 \lesssim_2 f_1$. Likewise $e_1 \not\lesssim_1 e_2$, $e_2 \not\lesssim_1 e_1$ and $f_1 \not\lesssim_2 f_2$, $f_2 \not\lesssim_2 f_1$.
4. $a_1 \lesssim_1 e_1$, $a_1 \neq e_1$ and $e_2 \neq \text{root}(s_1)$: Then we see that $a_2 \lesssim_2 f_1$, $a_2 \neq f_1$ and $f_2 \neq \text{root}(s_2)$ (again by the property of β_2 stated at the outset). Whereby $e_1 \lesssim_1 e_2$, $e_2 \lesssim_1 e_1$ and $f_1 \lesssim_2 f_2$, $f_2 \lesssim_2 f_1$. Likewise, $e_1 \not\lesssim_1 e_2$, $e_2 \not\lesssim_1 e_1$ and $f_1 \not\lesssim_2 f_2$, $f_2 \not\lesssim_2 f_1$.
5. $e_1 \neq a_1$, $e_1 \leq_1 a_1$: Then $f_1 \neq a_1$, $f_1 \leq_2 a_2$. Whereby $e_1 \leq_1 e_2$ and $f_1 \leq_2 f_2$. This is because $e_1 \leq_1 c_1$ and $f_1 \leq_2 c_2$ where c_1 and c_2 are resp. the parents of a_1 and a_2 in z_1 and z_2 . Also $e_1 \lesssim_1 e_2$, $e_2 \lesssim_1 e_1$ and $f_1 \lesssim_2 f_2$, $f_2 \lesssim_2 f_1$.
6. e_1 and e_2 are not related by \leq_1 or \lesssim_1 : Then f_1 and f_2 are also not related by \leq_2 or \lesssim_2 .

In all cases, we have that the pair (e_i, e_j) is in \leq_1 (resp. \lesssim_1) iff (f_i, f_j) is in \leq_2 (resp. \lesssim_2). \square

Proof of Theorem 10.2.2 for nested words

We first prove a composition lemma for nested words. Towards the statement of this lemma, we define the notion of *insert of a nested word v in a nested word u at a given position e of u* .

Definition 10.2.5 (Insert). Let $u = (A_u, \leq_u, \lambda_u, \rightsquigarrow_u)$ and $v = (A_v, \leq_v, \lambda_v, \rightsquigarrow_v)$ be given nested Σ -words, and let e be a position in u . The *insert of v in u at e* , denoted $u \uparrow_e v$, is a nested Σ -word defined as below.

1. If u and v have disjoint sets of positions, then $u \uparrow_e v = (A, \leq, \lambda, \rightsquigarrow)$ where
 - $A = A_u \sqcup A_v$
 - $\leq = \leq_u \cup \leq_v \cup \{(i, j) \mid i \in A_u, j \in A_v, i \leq_u e\} \cup \{(j, i) \mid i \in A_u, j \in A_v, e \leq_u i, e \neq i\}$
 - $\lambda(a) = \lambda_u(a)$ if $a \in A_u$, else $\lambda(a) = \lambda_v(a)$
 - $\rightsquigarrow = \rightsquigarrow_u \cup \rightsquigarrow_v$
2. If u and v have overlapping sets of positions, then let v_1 be an isomorphic copy of v whose set of positions is disjoint with that of u . Then $u \uparrow_e v$ is defined upto isomorphism as $u \uparrow_e v_1$.

In the special case that e is the last (under \leq_u) position of u , we denote $u \uparrow_e v$ as $u \cdot v$, and call the latter as the *concatenation of v with u* .

Lemma 10.2.6 (Composition lemma for nested words). *For a finite alphabet Σ , let $u_i, v_i \in \text{Nested-words}(\Sigma)$, and let e_i be a position in u_i for $i \in \{1, 2\}$. Then the following hold for each $m \in \mathbb{N}$.*

1. *If $(u_1, e_1) \equiv_{m, \mathcal{L}} (u_2, e_2)$ and $v_1 \equiv_{m, \mathcal{L}} v_2$, then $(u_1 \uparrow_{e_1} v_1) \equiv_{m, \mathcal{L}} (u_2 \uparrow_{e_2} v_2)$.*
2. *$u_1 \equiv_{m, \mathcal{L}} u_2$ and $v_1 \equiv_{m, \mathcal{L}} v_2$, then $u_1 \cdot v_1 \equiv_{m, \mathcal{L}} u_2 \cdot v_2$.*

Proof. We give the proof for $\mathcal{L} = \text{MSO}$. The proof for $\mathcal{L} = \text{FO}$ is similar.

The winning strategy S for the duplicator in the m -round MSO-EF game between $u_1 \uparrow_{e_1} v_1$ and $u_2 \uparrow_{e_2} v_2$ is simply the composition of the winning strategies S_1 , resp. S_2 , of the duplicator in the m -round MSO-EF game between (u_1, e_1) and (u_2, e_2) , resp. v_1 and v_2 . Formally, S is defined as follows.

1. Point move: If the spoiler picks an element of u_1 , resp. v_1 , from $u_1 \uparrow_{e_1} v_1$, then the duplicator picks the element of u_2 , resp. v_2 , from $u_2 \uparrow_{e_2} v_2$, that is given by the strategy S_1 , resp. S_2 . A similar choice of an element from $u_1 \uparrow_{e_1} v_1$ is made by the duplicator if the spoiler picks an element from $u_2 \uparrow_{e_2} v_2$.

2. Set move: If the spoiler picks a set Z from $u_1 \uparrow_{e_1} v_1$, then let $Z = X \sqcup Y$ where X is a subset of positions of u_1 and Y is a subset of positions of v_1 . Then the duplicator picks the set Z' from $u_2 \uparrow_{e_2} v_2$ where $Z' = X' \sqcup Y'$, X' is the subset of positions of u_2 that is chosen by the duplicator in response to X according to strategy S_1 , and Y' is the subset of positions of v_2 that is chosen by the duplicator in response to Y according to strategy S_2 . A similar choice of a set from $u_1 \uparrow_{e_1} v_1$ is made by the duplicator if the spoiler picks a set from $u_2 \uparrow_{e_2} v_2$.

It is easy to see that S is a winning strategy in the MSO-EF game between $u_1 \uparrow_{e_1} v_1$ and $u_2 \uparrow_{e_2} v_2$. \square

Towards the proof of Theorem 10.2.2 for nested words, we first observe that each nested Σ -word has a natural representation using a representation-feasible tree over $\Sigma_{\text{int}} \cup \Sigma_{\text{leaf}}$, where $\Sigma_{\text{leaf}} = \Sigma \cup (\Sigma \times \Sigma)$, and $\Sigma_{\text{int}} = \Sigma_{\text{leaf}} \cup \{\circ\}$. We demonstrate this for the example of the nested Σ -word $w = (abaabba, \{(2, 6), (4, 5)\})$, where $\Sigma = \{a, b\}$. See Figure 10.1.





Figure 10.1: Nested Σ -word as a tree
over $\Sigma \cup (\Sigma \times \Sigma) \cup \{\circ\}$



$\text{Str}(t_{\geq y}) = (a, \emptyset)$
 $\text{Str}(t_{\geq z}) = (ab, \{(1, 2)\})$
 $\text{Str}(t_{\geq w}) = (baabb, \{(1, 5), (3, 4)\})$
 $\text{Str}(t_{\geq u}) = (abaabba, \{(2, 6), (4, 5)\}) = w$

Figure 10.2: Tree over $\Sigma \cup (\Sigma \times \Sigma) \cup \{\circ\}$
as a nested Σ -word

Formally, each non-empty nested Σ -word can be seen to be of one of two types. A non-empty nested word $u = (A, \leq, \lambda, \rightsquigarrow)$ is said to be of *type A* if either \rightsquigarrow is empty and $|A| = 1$ (i.e. u is really a Σ -word of length 1), or for the minimum and maximum (under \leq) positions i and j respectively of u , it is the case that $i \rightsquigarrow j$. A non-empty nested word u is said to be of *type B* if it is not of type A. It is easy to see that a type B nested word can be written as a concatenation

of type A nested words. We describe inductively, the tree-representation of $u = (A, \leq, \lambda, \rightsquigarrow)$ below. We have three cases.

1. u is empty: Then the tree t over $\Sigma_{\text{int}} \cup \Sigma_{\text{leaf}}$ representing u is the empty tree.
2. u is of type A: If \rightsquigarrow is empty and $|A| = 1$, then let the only element of A be labeled (by λ) with the letter a , where $a \in \Sigma$. Then the tree t over $\Sigma_{\text{int}} \cup \Sigma_{\text{leaf}}$ representing u is a singleton whose only node is labeled with a .

Else, let u_1 be the nested sub- Σ -word of u induced by the positions $l \in A$ such that $i \leq l \leq j, l \neq i$ and $l \neq j$, where i and j are respectively the minimum and maximum (under \leq) positions of u . Let t_1 be the tree over $\Sigma_{\text{int}} \cup \Sigma_{\text{leaf}}$ representing u_1 , if the latter is not empty. Then the tree t over $\Sigma_{\text{int}} \cup \Sigma_{\text{leaf}}$ representing u is defined as follows. If u_1 is empty, then t is a singleton whose only node is labeled with the label $(\lambda(i), \lambda(j))$. Else, t is such that (i) the label of the $\text{root}(t)$ is $(\lambda(i), \lambda(j))$ and (ii) the only child of $\text{root}(t)$ is $\text{root}(t_1)$, i.e. t_1 is the only child subtree of $\text{root}(t)$.

3. u is of type B: Then u can be written as $u = u_1 \cdots u_n$ where u_i is a type A nested Σ -word for $i \in \{1, \dots, n\}$. Let t_i be the tree over $\Sigma_{\text{int}} \cup \Sigma_{\text{leaf}}$ representing u_i . Then the tree t over $\Sigma_{\text{int}} \cup \Sigma_{\text{leaf}}$ representing u is such that (i) the label of $\text{root}(t)$ is \circ , and (ii) the children of $\text{root}(t)$ in “increasing order” are $\text{root}(t_1), \dots, \text{root}(t_n)$.

Conversely, each representation-feasible tree t over $\Sigma_{\text{int}} \cup \Sigma_{\text{leaf}}$ represents a nested Σ -word u_t .

We demonstrate this for the example of the nested Σ -word $w = (\text{abaabba}, \{(2, 6), (4, 5)\})$ where $\Sigma = \{a, b\}$, in Figure 10.2. Formally, we see this inductively as follows.

1. If t is empty, then u_t is the empty nested Σ -word.
2. If $t = (O, \lambda)$ contains only a single node, say e , then there are two cases. If $\lambda(e) \in \Sigma$, then $u_t = (A_t, \leq_t, \lambda_t, \rightsquigarrow_t)$ where $A_t = \{e_1\}$, $\leq_t = \{(e_1, e_1)\}$, $\lambda_t(e_1) = \lambda(e)$ and $\rightsquigarrow_t = \emptyset$. Else, i.e. if $\lambda(e) = (a, b) \in \Sigma \times \Sigma$, then $u_t = (A_t, \leq_t, \lambda_t, \rightsquigarrow_t)$ where $A_t = \{e_1, e_2\}$, $\leq_t = \{(e_1, e_1), (e_1, e_2), (e_2, e_2)\}$, $\lambda_t(e_1) = a$, $\lambda_t(e_2) = b$ and $\rightsquigarrow_t = \{(e_1, e_2)\}$.
3. If $t = (O, \lambda)$ contains more than one node, then let t_1, \dots, t_n be, in “increasing order”, the subtrees of t rooted at the children of $\text{root}(t)$. Let $v = u_{t_1} \cdots u_{t_n}$, where u_{t_i} is the nested Σ -word represented by t_i , for $i \in \{1, \dots, n\}$. If $\lambda(\text{root}(t)) = \circ$, then $u_t = v$. Else, suppose that $\lambda(\text{root}(t)) = (a, b)$ where $a, b \in \Sigma$. Let w be the 2-letter Σ -word given by $w = (A_w, \leq_w, \lambda_w, \rightsquigarrow_w)$ where $A_w = \{e_1, e_2\}$, $\leq_w = \{(e_1, e_1), (e_1, e_2), (e_2, e_2)\}$, $\lambda_w(e_1) = a$, $\lambda_w(e_2) = b$ and $\rightsquigarrow_w = \{(e_1, e_2)\}$. Then $u_t = w \uparrow_{e_1} v$.

We now prove Theorem 10.2.2 for nested words. It suffices to show that $\mathcal{L}\text{-EBSP}(\cdot, k)$ holds

with a computable witness function for $\text{Nested-words}(\Sigma)$. That any regular subclass of $\text{Nested-words}(\Sigma)$ satisfies $\mathcal{L}\text{-EBSP}(\cdot, k)$ follows, because (i) a regular subclass of $\text{Nested-words}(\Sigma)$ is MSO definable, (ii) MSO-EBSP(\cdot, k) and the computability of witness function are preserved under MSO definable subclasses (Lemma 10.4.1(4)), and (iii) MSO-EBSP(\cdot, k) implies FO-EBSP(\cdot, k), and any witness function for the former is also a witness function for the latter (Lemma 9.3.1(4)).

Let $\mathcal{S} = \text{Nested-words}(\Sigma)$. Consider the class \mathcal{S}_p where $p \geq 1$. Let \mathcal{T} be the class of all representation-feasible trees over $\Sigma'_{\text{int}} \cup \Sigma'_{\text{leaf}}$ where $\Sigma'_{\text{int}} = \Sigma'_{\text{leaf}} \cup \{\circ\}$, $\Sigma'_{\text{leaf}} = \Sigma' \cup (\Sigma' \times \Sigma')$ and $\Sigma' = \Sigma \times \{0, \dots, p-1\}$. Let $\text{Str} : \mathcal{T} \rightarrow \mathcal{S}_p$ be the map given by $\text{Str}(t)$ is the nested Σ' -word represented by t as described above. That Str is size effective and onto upto isomorphism, is clear. That \mathcal{T} is closed under subtrees, and that Str satisfies the conditions A.1.a, B.1 and B.2 of Section 10.1 are easy to see. That Str satisfies A.1.b for $m \geq 0$ follows directly from Lemma 10.2.6. We show below that Str satisfies A.2 for all $m \geq 2$. Then Str is \mathcal{L} -height-reduction favourable and \mathcal{L} -degree-reduction favourable, whereby \mathcal{T} and Str are of \mathcal{L} -RFSE type II. Then \mathcal{S} admits an \mathcal{L} -RFSE tree representation schema, whereby using Lemma 10.1.2, it follows that $\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ holds with a computable witness function.

Let $t, s_1, s_2 \in \mathcal{T}$ be trees of size ≥ 2 such that the roots of all these trees have the same label, and suppose $z_i = t \odot s_i \in \mathcal{T}$ for $i \in \{1, 2\}$. Assume $\text{Str}(s_1) \equiv_{m, \mathcal{L}} \text{Str}(s_2)$ for $m \geq 2$. If the label of $\text{root}(t)$ is \circ , then $\text{Str}(z_i) = \text{Str}(t) \cdot \text{Str}(s_i)$ for $i \in \{1, 2\}$ whereby from Lemma 10.2.6, $\text{Str}(z_1) \equiv_{m, \mathcal{L}} \text{Str}(z_2)$. Else suppose the label of $\text{root}(t)$ is (a, b) where $a, b \in \Sigma$. Let u be the concatenation, in “increasing order”, of the nested Σ -words represented by the subtrees of t rooted at the children of the root of t . Likewise, let v_i be the concatenation, in “increasing order”, of the nested Σ -words represented by the subtrees of s_i rooted at the children of the root of s_i , for $i \in \{1, 2\}$. Let w be the 2-letter Σ -word given by $w = (A_w, \leq_w, \lambda_w, \rightsquigarrow_w)$ where $A_w = \{e_1, e_2\}$, $\leq_w = \{(e_1, e_1), (e_1, e_2), (e_2, e_2)\}$, $\lambda_w(e_1) = a$, $\lambda_w(e_2) = b$ and $\rightsquigarrow_w = \{(e_1, e_2)\}$. Then $\text{Str}(z_i) = w \uparrow_{e_1} (u \cdot v_i)$. Now observe that since $\text{Str}(s_1) \equiv_{m, \mathcal{L}} \text{Str}(s_2)$, we have $v_1 \equiv_{m, \mathcal{L}} v_2$ for each $m \geq 2$. Then by Lemma 10.2.6, we have $\text{Str}(z_1) \equiv_{m, \mathcal{L}} \text{Str}(z_2)$, completing the proof.

10.3 n -partite cographs

The class of n -partite cographs was introduced by Ganian et. al. in [31]. An n -partite cograph G is a graph that admits an n -partite *cotree representation* t . Here t is an unordered tree whose

leaves are exactly the vertices of G , and are labeled with labels from $[n] = \{1, \dots, n\}$. Each internal node v of t is labeled with a binary symmetric function $f_v : [n] \times [n] \rightarrow \{0, 1\}$ such that two vertices a and b of G with respective labels i and j , are adjacent in G iff the greatest common ancestor of a and b in t , call it c , is such that $f_c(i, j) = 1$. Given below is an example of an n -partite cograph G and a cotree representation t of it.




Figure 10.3: n -partite cograph G and an n -partite cotree representation t of G

Given a finite alphabet Σ , a Σ -labeled n -partite cograph is a pair (G, ν) where G is an n -partite cograph and $\nu : V \rightarrow \Sigma$ is a labeling function. Recall that given a class \mathcal{S}_1 of structures and a subclass \mathcal{S}_2 of \mathcal{S}_1 , we say \mathcal{S}_2 is *hereditary over \mathcal{S}_1* , if \mathcal{S}_2 is *PS* over \mathcal{S}_1 (see the last paragraph of Chapter 7). A class \mathcal{S} is *hereditary* if it is hereditary over the class of all (finite) structures.

The central result of this section can now be stated as follows.

Theorem 10.3.1. *Given $n, k \in \mathbb{N}$, let $\text{Labeled-}n\text{-partite-cographs}(\Sigma)$ be the class of all Σ -labeled n -partite cographs. Let \mathcal{S} be any subclass of $\text{Labeled-}n\text{-partite-cographs}(\Sigma)$, that is hereditary over the latter. Then $\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ holds with a computable witness function. Consequently, each of the following classes of graphs satisfies $\mathcal{L}\text{-EBSP}(\cdot, k)$ with a computable witness function for each $k \geq 0$.*

1. Any hereditary class of n -partite cographs, for each $n \in \mathbb{N}$.
2. Any hereditary class of graphs of bounded shrub-depth.
3. Any hereditary class of graphs of bounded SC-depth.
4. Any hereditary class of graphs of bounded tree-depth.
5. Any hereditary class of cographs.

Proof. We first show that the class $\mathcal{S} = n\text{-partite-cographs}$, where $n\text{-partite-cographs}$ is the class of all n -partite cographs, admits an MSO-RFSE tree representation schema. Then by

Lemma 10.1.2, we have $\text{MSO-EBSP}(\mathcal{S}_p, k)$ holds with a computable witness function for each $p, k \in \mathbb{N}$. Then $\mathcal{L}\text{-EBSP}(\mathcal{S}_p, k)$ holds with a computable witness function for each $p, k \in \mathbb{N}$ by Lemma 9.3.1(4). Since there is a 1-1 correspondence between $\text{Labeled-}n\text{-partite-cographs}(\Sigma)$ and \mathcal{S}_p if $|\Sigma| = p$, it follows that $\mathcal{L}\text{-EBSP}(\cdot, k)$ holds with a computable witness function for $\text{Labeled-}n\text{-partite-cographs}(\Sigma)$. Whereby, the same holds of any subclass of $\text{Labeled-}n\text{-partite-cographs}(\Sigma)$ that is hereditary over the latter by Lemma 10.4.1(1). That $\mathcal{L}\text{-EBSP}(\cdot, k)$ holds with a computable witness function for the various specific classes mentioned in the statement of this result, follows from the fact that these classes are hereditary subclasses of n -partite-cographs, and the fact that n -partite-cographs is itself hereditary [31].

Consider \mathcal{S}_p for $p \geq 1$. Let $\Sigma_{\text{leaf}} = [n] \times \{0, \dots, p-1\}$ and $\Sigma_{\text{int}} = \{f \mid f : [n] \times [n] \rightarrow \{0, 1\}\}$. Let \mathcal{T} be the class of all representation-feasible $(\Sigma_{\text{int}} \cup \Sigma_{\text{leaf}})$ -trees. Then \mathcal{T} is closed under subtrees. Let $\text{Str} : \mathcal{T} \rightarrow \mathcal{S}_p$ be such that for $t = (O, \lambda) \in \mathcal{T}$, we have $\text{Str}(t) = (G, \nu)$ where (i) G is the n -partite cograph represented by the unordered tree obtained from t by “forgetting” the ordering among the children of t and by dropping the second component of the labels of the leaves of t , and (ii) ν is such that for any vertex a of G (which is a leaf node of t), it is the case that $\nu(a)$ is the second component of $\lambda(a)$. It is easily seen that Str is size effective and onto upto isomorphism. It is also easy to see that Str satisfies conditions A.1.a, B.1 and B.2 of Section 10.1. We now show below that Str satisfies for $\mathcal{L} = \text{MSO}$, the conditions A.1.b and A.2 for $m \geq 0$; then \mathcal{T} and Str are of $\mathcal{L}\text{-RFSE-type II}$. Whereby, \mathcal{S} admits an $\mathcal{L}\text{-RFSE}$ tree representation schema, completing the proof.

For our proof, we need the following composition lemma. We prove this lemma towards the end of this section.

Lemma 10.3.2 (Composition lemma for n -partite cographs). *For $i \in \{1, 2\}$, let $(G_i, \nu_{i,1})$ and $(H_i, \nu_{i,2})$ be graphs in \mathcal{S}_p . Suppose t_i and s_i are trees of \mathcal{T} such that $\text{Str}(t_i) = (G_i, \nu_{i,1})$, $\text{Str}(s_i) = (H_i, \nu_{i,2})$, and the labels of $\text{root}(t_i)$ and $\text{root}(s_i)$ are the same. Let $z_i = t_i \odot s_i$ and $\text{Str}(z_i) = (Z_i, \nu_i)$ for $i \in \{1, 2\}$. For each $m \in \mathbb{N}$, if $(G_1, \nu_{1,1}) \equiv_{m, \text{MSO}} (G_2, \nu_{2,1})$ and $(H_1, \nu_{1,2}) \equiv_{m, \text{MSO}} (H_2, \nu_{2,2})$, then $(Z_1, \nu_1) \equiv_{m, \text{MSO}} (Z_2, \nu_2)$.*

We now show that Lemma 10.3.2 implies that Str satisfies A.1.b and A.2 for $m \geq 0$ and $\mathcal{L} = \text{MSO}$.

That Str satisfies A.2 for $m \geq 0$ and $\mathcal{L} = \text{MSO}$ follows easily from Lemma 10.3.2. Let $t, s_1, s_2 \in \mathcal{T}$ be trees of size ≥ 2 such that the roots of all these trees have the same label. For

$i \in \{1, 2\}$, suppose $z_i = t \odot s_i \in \mathcal{T}$ and that $\text{Str}(s_1) \equiv_{m, \text{MSO}} \text{Str}(s_2)$. Let $t_i = t$, $(G_i, \nu_{i,1}) = \text{Str}(t_i)$ and $(H_i, \nu_{i,2}) = \text{Str}(s_i)$ for $i \in \{1, 2\}$. It now follows directly from Lemma 10.3.2, that $\text{Str}(z_1) \equiv_{m, \text{MSO}} \text{Str}(z_2)$.

To see that Str satisfies A.1.b for $m \geq 0$ and $\mathcal{L} = \text{MSO}$, let $t, s_1 \in \mathcal{T}$ and a be a child of $\text{root}(t)$. Suppose that $s_2 = t_{\geq a}$ and $z = t[s_2 \mapsto s_1] \in \mathcal{T}$, and that $\text{Str}(s_1) \equiv_{m, \text{MSO}} \text{Str}(s_2)$. For $i \in \{1, 2\}$, let s'_i be the tree in \mathcal{T} obtained by making the root of s_i , the sole child of a new node whose label is the same as the label of $\text{root}(t)$ in t . Using the notation introduced in Section 10.2, if s_3 is the singleton tree whose sole node, say b , is labeled with the same label as that of $\text{root}(t)$ in t , then $s'_i = s_3 \cdot_b^\uparrow s_i$. It is easy to verify that $\text{Str}(s_i) = \text{Str}(s'_i)$, whereby $\text{Str}(s'_1) \equiv_{m, \text{MSO}} \text{Str}(s'_2)$. Let y_1 , resp. y_2 , be the subtree of t obtained by deleting the subtrees of t rooted at the children of $\text{root}(t)$ that are “greater than or equal to” a , resp. “less than or equal to” a , under the ordering of the children of $\text{root}(t)$ in t . Then $t = (y_1 \odot s'_2) \odot y_2$ and $z = (y_1 \odot s'_1) \odot y_2$. Since $\text{Str}(s'_1) \equiv_{m, \text{MSO}} \text{Str}(s'_2)$, we have by Lemma 10.3.2 that $\text{Str}(y_1 \odot s'_1) \equiv_{m, \text{MSO}} \text{Str}(y_1 \odot s'_2)$, whereby $\text{Str}(t) \equiv_{m, \text{MSO}} \text{Str}(z)$, showing that Str satisfies A.1.b for $m \geq 0$ and $\mathcal{L} = \text{MSO}$, completing the proof. \square

Proof of Lemma 10.3.2. We can assume w.l.o.g. that t_i and s_i have disjoint sets of nodes for $i \in \{1, 2\}$. Let the set of vertices of $\text{Str}(t_i)$ and $\text{Str}(s_i)$ be $V\text{-Str}(t_i)$ and $V\text{-Str}(s_i)$ respectively. Then the vertex set $V\text{-Str}(z_i)$ of $\text{Str}(z_i)$ is $V\text{-Str}(t_i) \sqcup V\text{-Str}(s_i)$ for $i \in \{1, 2\}$.

Let \mathbf{S}_t , resp. \mathbf{S}_s , be the strategy of the duplicator in the m -round MSO-EF game between $\text{Str}(t_1)$ and $\text{Str}(t_2)$, resp. between $\text{Str}(s_1)$ and $\text{Str}(s_2)$. For the m -round MSO-EF game between $\text{Str}(z_1)$ and $\text{Str}(z_2)$, the duplicator follows the following strategy, call it \mathbf{R} .

- Point move: If the spoiler chooses a vertex from $V\text{-Str}(t_1)$ (resp. $V\text{-Str}(t_2)$), then the duplicator chooses a vertex from $V\text{-Str}(t_2)$ (resp. $V\text{-Str}(t_1)$) according to \mathbf{S}_t . Else, if the spoiler chooses a vertex from $V\text{-Str}(s_1)$ (resp. $V\text{-Str}(s_2)$), then the duplicator chooses a vertex from $V\text{-Str}(s_2)$ (resp. $V\text{-Str}(s_1)$) according to \mathbf{S}_s .
- Set move: If the spoiler chooses a set, say U , from $V\text{-Str}(z_1)$ (resp. $V\text{-Str}(z_2)$), then let $X = U \cap V\text{-Str}(t_1)$ (resp. $X = U \cap V\text{-Str}(t_2)$) and $Y = U \cap V\text{-Str}(s_1)$ (resp. $Y = U \cap V\text{-Str}(s_2)$). Let X' be the subset of $V\text{-Str}(t_2)$ (resp. $V\text{-Str}(t_1)$) that is picked according to the strategy \mathbf{S}_t in response to the choice of X in $V\text{-Str}(t_1)$ (resp. $V\text{-Str}(t_2)$). Likewise, let Y' be the subset of $V\text{-Str}(s_2)$ (resp. $V\text{-Str}(s_1)$) that is picked according to \mathbf{S}_s in response to the choice of Y in $V\text{-Str}(s_1)$ (resp. $V\text{-Str}(s_2)$). Then the set U' picked by the duplicator from $V\text{-Str}(z_2)$ according to strategy \mathbf{R} is given by $U' = X' \sqcup Y'$.

We now show that \mathbf{R} is a winning strategy for the duplicator.

Let at the end of m rounds, the vertices and sets chosen from $\text{Str}(z_1)$, resp. $\text{Str}(z_2)$, be a_1, \dots, a_p and A_1, \dots, A_r , resp. b_1, \dots, b_p and B_1, \dots, B_r , where $p + r = m$. Let $A_l^1 = A_l \cap \text{V-Str}(t_1)$, $A_l^2 = A_l \cap \text{V-Str}(s_1)$, $B_l^1 = B_l \cap \text{V-Str}(t_2)$ and $B_l^2 = B_l \cap \text{V-Str}(s_2)$ for $l \in \{1, \dots, r\}$.

It is easy to see that the labels of a_i and b_i are the same for all $i \in \{1, \dots, p\}$. Also by the description of \mathbf{R} given above it is easy to check for all $i \in \{1, \dots, p\}$ that $a_i \in \text{V-Str}(t_1)$ iff $b_i \in \text{V-Str}(t_2)$ and $a_i \in \text{V-Str}(s_1)$ iff $b_i \in \text{V-Str}(s_2)$. Likewise, for all $l \in \{1, \dots, r\}$ and $i \in \{1, \dots, p\}$, we have $a_i \in A_l^1$ iff $b_i \in B_l^1$ and $a_i \in A_l^2$ iff $b_i \in B_l^2$, whereby $a_i \in A_l$ iff $b_i \in B_l$.

Consider a_i, a_j for $i \neq j$ and $i, j \in \{1, \dots, p\}$. We show below that a_i, a_j are adjacent in $\text{Str}(z_1)$ iff b_i, b_j are adjacent in $\text{Str}(z_2)$. This would show that $a_i \mapsto b_i$ is a partial isomorphism between $(\text{Str}(z_1), A_1, \dots, A_r)$ and $(\text{Str}(z_2), B_1, \dots, B_r)$ completing the proof. We have the following three cases:

1. Each of a_i and a_j is from $\text{V-Str}(t_1)$: Then by the description of \mathbf{R} above, we have that (i) b_i and b_j are both from $\text{V-Str}(t_2)$ and (ii) a_i, a_j are adjacent in $\text{Str}(t_1)$ iff b_i, b_j are adjacent in $\text{Str}(t_2)$. Observe that a_i, a_j are adjacent in $\text{Str}(t_1)$ iff a_i, a_j are adjacent in $\text{Str}(z_1)$. Likewise, b_i, b_j are adjacent in $\text{Str}(t_2)$ iff b_i, b_j are adjacent in $\text{Str}(z_2)$. Then a_i, a_j are adjacent in $\text{Str}(z_1)$ iff b_i, b_j are adjacent in $\text{Str}(z_2)$.
2. Each of a_i and a_j is from $\text{V-Str}(s_1)$: Reasoning similarly as in the previous case, we can show that a_i, a_j are adjacent in $\text{Str}(z_1)$ iff b_i, b_j are adjacent in $\text{Str}(z_2)$.
3. W.l.o.g. $a_i \in \text{V-Str}(t_1)$ and $a_j \in \text{V-Str}(s_1)$: Then $b_i \in \text{V-Str}(t_2)$ and $b_j \in \text{V-Str}(s_2)$. Observe now that the greatest common ancestor of a_i and a_j in z_1 is $\text{root}(z_1)$, and the greatest common ancestor of b_i and b_j in z_2 is $\text{root}(z_2)$. Since (i) the labels of $\text{root}(z_1)$ and $\text{root}(z_2)$ are the same (by assumption) and (ii) the label of a_i (resp. a_j) in z_1 = label of a_i (resp. a_j) in $\text{Str}(z_1)$ = label of b_i (resp. b_j) in $\text{Str}(z_2)$ = label of b_i (resp. b_j) in z_2 , it follows by the definition of an n -partite cograph that a_i, a_j are adjacent in $\text{Str}(z_1)$ iff b_i, b_j are adjacent in $\text{Str}(z_2)$.

□

10.4 Closure properties of \mathcal{L} -EBSP(\cdot, \cdot)

10.4.1 Closure of \mathcal{L} -EBSP(\cdot, \cdot) under set-theoretic operations

Lemma 10.4.1. *Given classes \mathcal{S}_1 and \mathcal{S}_2 of finite structures, and $k_1, k_2 \in \mathbb{N}$, suppose that \mathcal{L} -EBSP(\mathcal{S}_i, k_i) is true for each $i \in \{1, 2\}$, and suppose $\theta_{(\mathcal{S}_i, k_i, \mathcal{L})}$ is a witness function for \mathcal{L} -EBSP(\mathcal{S}_i, k_i). Then the following hold.*

1. *If \mathcal{S} is any subclass of \mathcal{S}_i that is hereditary over \mathcal{S}_i , where $i \in \{1, 2\}$, then \mathcal{L} -EBSP(\mathcal{S}, k) is true for $k = k_i$, with witness function $\theta_{(\mathcal{S}, k, \mathcal{L})}$ given by $\theta_{(\mathcal{S}, k, \mathcal{L})} = \theta_{(\mathcal{S}_i, k_i, \mathcal{L})}$.*
2. *If $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2$, then \mathcal{L} -EBSP(\mathcal{S}, k) is true for $k = \min(k_1, k_2)$, with witness function $\theta_{(\mathcal{S}, k, \mathcal{L})}$ given by $\theta_{(\mathcal{S}, k, \mathcal{L})} = \max(\theta_{(\mathcal{S}_1, k_1, \mathcal{L})}, \theta_{(\mathcal{S}_2, k_2, \mathcal{L})})$.*
3. *If $\mathcal{S} = \mathcal{S}_1 \cap \mathcal{S}_2$ and \mathcal{S}_1 is hereditary, then \mathcal{L} -EBSP(\mathcal{S}, k) is true for $k = k_2$, with witness function $\theta_{(\mathcal{S}, k, \mathcal{L})}$ given by $\theta_{(\mathcal{S}, k, \mathcal{L})} = \theta_{(\mathcal{S}_2, k_2, \mathcal{L})}$. If \mathcal{S}_2 is also hereditary, then \mathcal{L} -EBSP(\mathcal{S}, k) is true for $k = \max(k_1, k_2)$, with witness function $\theta_{(\mathcal{S}, k, \mathcal{L})}$ given by $\theta_{(\mathcal{S}, k, \mathcal{L})} = \max(\theta_{(\mathcal{S}_1, k_1, \mathcal{L})}, \theta_{(\mathcal{S}_2, k_2, \mathcal{L})})$.*
4. *If \mathcal{S} is a subclass of \mathcal{S}_i that is definable over \mathcal{S}_i by an \mathcal{L} sentence of rank r , then \mathcal{L} -EBSP(\mathcal{S}, k) is true for $k = k_i$, with witness function $\theta_{(\mathcal{S}, k, \mathcal{L})}$ given by $\theta_{(\mathcal{S}, k, \mathcal{L})}(m) = \theta_{(\mathcal{S}_i, k_i, \mathcal{L})}(r)$ if $m \leq r$, else $\theta_{(\mathcal{S}, k, \mathcal{L})}(m) = \theta_{(\mathcal{S}_i, k_i, \mathcal{L})}(m)$. It follows that for \mathcal{S} as aforementioned, if $\overline{\mathcal{S}}$ is the complement of \mathcal{S} in \mathcal{S}_i , then \mathcal{L} -EBSP($\overline{\mathcal{S}}, k$) is also true for $k = k_i$, with witness function $\theta_{(\overline{\mathcal{S}}, k, \mathcal{L})}$ given by $\theta_{(\overline{\mathcal{S}}, k, \mathcal{L})} = \theta_{(\mathcal{S}, k, \mathcal{L})}$ where $\theta_{(\mathcal{S}, k, \mathcal{L})}$ is as aforementioned.*

Proof. Let $m \in \mathbb{N}$ be given.

Part 1: Consider $\mathfrak{A} \in \mathcal{S}$ and let \bar{a} be a k -tuple from \mathfrak{A} where $k = k_i$. Since \mathcal{L} -EBSP(\mathcal{S}_i, k_i) is true, there exists $\mathfrak{B} \in \mathcal{S}_i$ such that \mathcal{L} -EBSP-condition($\mathcal{S}_i, \mathfrak{A}, \mathfrak{B}, k_i, m, \bar{a}, \theta_{(\mathcal{S}_i, k_i, \mathcal{L})}$) is true. Then since $\mathfrak{B} \subseteq \mathfrak{A}$ and \mathcal{S} is hereditary over \mathcal{S}_i , we have $\mathfrak{B} \in \mathcal{S}$; whence \mathcal{L} -EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, k, m, \bar{a}, \theta_{(\mathcal{S}, k, \mathcal{L})}$) is true, where $\theta_{(\mathcal{S}, k, \mathcal{L})} = \theta_{(\mathcal{S}_i, k_i, \mathcal{L})}$.

Part 2: Consider $\mathfrak{A} \in \mathcal{S}$ and let \bar{a} be a k -tuple from \mathfrak{A} where $k = \min(k_1, k_2)$. Since $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2$, assume w.l.o.g. that $\mathfrak{A} \in \mathcal{S}_1$. Let b be an element of \bar{a} and let \bar{a}_1 be a k_1 -tuple whose first k components form exactly the tuple \bar{a} and in which b is the element at all the indices $k+1, \dots, k_1$. Since \mathcal{L} -EBSP(\mathcal{S}_1, k_1) is true, there exists $\mathfrak{B} \in \mathcal{S}_1$ such that \mathcal{L} -EBSP-condition($\mathcal{S}_1, \mathfrak{A}, \mathfrak{B}, k_1, m, \bar{a}_1, \theta_{(\mathcal{S}_1, k_1, \mathcal{L})}$) is true. Then $\mathfrak{B} \in \mathcal{S}$. Further since $\text{tp}_{\mathfrak{B}, \bar{a}_1, m, \mathcal{L}}(\bar{x}) = \text{tp}_{\mathfrak{A}, \bar{a}_1, m, \mathcal{L}}(\bar{x})$, it follows that $\text{tp}_{\mathfrak{B}, \bar{a}, m, \mathcal{L}}(\bar{x}) = \text{tp}_{\mathfrak{A}, \bar{a}, m, \mathcal{L}}(\bar{x})$; whence \mathcal{L} -EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, k, m, \bar{a}, \theta_{(\mathcal{S}, k, \mathcal{L})}$) is true, where $\theta_{(\mathcal{S}, k, \mathcal{L})} = \max(\theta_{(\mathcal{S}_1, k_1, \mathcal{L})}, \theta_{(\mathcal{S}_2, k_2, \mathcal{L})})$.

Part 3: Consider $\mathfrak{A} \in \mathcal{S}$ and let \bar{a} be a k -tuple from \mathfrak{A} where $k = k_2$. Since \mathcal{L} -EBSP(\mathcal{S}_2, k_2)

is true, there exists $\mathfrak{B} \in \mathcal{S}_2$ such that \mathcal{L} -EBSP-condition($\mathcal{S}_2, \mathfrak{A}, \mathfrak{B}, k_2, m, \bar{a}, \theta_{(\mathcal{S}_2, k_2, \mathcal{L})}$) is true. Since $\mathfrak{B} \subseteq \mathfrak{A}, \mathfrak{A} \in \mathcal{S}_1$ and \mathcal{S}_1 is hereditary, we have $\mathfrak{B} \in \mathcal{S}_1$, and hence $\mathfrak{B} \in \mathcal{S}$. Then \mathcal{L} -EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, k, m, \bar{a}, \theta_{(\mathcal{S}, k, \mathcal{L})}$) is true, where $\theta_{(\mathcal{S}, k, \mathcal{L})} = \theta_{(\mathcal{S}_2, k_2, \mathcal{L})}$. If \mathcal{S}_2 is also hereditary, then let \bar{a} be a k -tuple from \mathfrak{A} where $k = \max(k_1, k_2)$. W.l.o.g., suppose $k_1 \geq k_2$, so that $k = k_1$. Since \mathcal{L} -EBSP(\mathcal{S}_1, k_1) is true, there exists $\mathfrak{B} \in \mathcal{S}_1$ such that \mathcal{L} -EBSP-condition($\mathcal{S}_1, \mathfrak{A}, \mathfrak{B}, k_1, m, \bar{a}, \theta_{(\mathcal{S}_1, k_1, \mathcal{L})}$) is true. Since $\mathfrak{B} \subseteq \mathfrak{A}, \mathfrak{A} \in \mathcal{S}_2$ and \mathcal{S}_2 is hereditary, we have $\mathfrak{B} \in \mathcal{S}_2$, and hence $\mathfrak{B} \in \mathcal{S}$. Then \mathcal{L} -EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, k, m, \bar{a}, \theta_{(\mathcal{S}, k, \mathcal{L})}$) is true, where $\theta_{(\mathcal{S}, k, \mathcal{L})} = \max(\theta_{(\mathcal{S}_1, k_1, \mathcal{L})}, \theta_{(\mathcal{S}_2, k_2, \mathcal{L})})$.

Part 4: Consider $\mathfrak{A} \in \mathcal{S}$ and let \bar{a} be a k -tuple from \mathfrak{A} where $k = k_i$. Let $\theta_{(\mathcal{S}, k, \mathcal{L})}$ and $\theta_{(\bar{\mathcal{S}}, k, \mathcal{L})}$ be functions as defined in the statement of this part. Since $\theta_{(\mathcal{S}_i, k_i, \mathcal{L})}$ is monotonic (see Remark 9.2), so are $\theta_{(\mathcal{S}, k, \mathcal{L})}$ and $\theta_{(\bar{\mathcal{S}}, k, \mathcal{L})}$.

Suppose $m \leq r$. Since \mathcal{L} -EBSP(\mathcal{S}_i, k) is true, there exists $\mathfrak{B} \in \mathcal{S}_i$ such that \mathcal{L} -EBSP-condition($\mathcal{S}_i, \mathfrak{A}, \mathfrak{B}, k, r, \bar{a}, \theta_{(\mathcal{S}_i, k, \mathcal{L})}$) is true. Then since $\text{tp}_{\mathfrak{B}, \bar{a}, r, \mathcal{L}}(\bar{x}) = \text{tp}_{\mathfrak{A}, \bar{a}, r, \mathcal{L}}(\bar{x})$, we have (i) $\text{tp}_{\mathfrak{B}, \bar{a}, m, \mathcal{L}}(\bar{x}) = \text{tp}_{\mathfrak{A}, \bar{a}, m, \mathcal{L}}(\bar{x})$ since $m \leq r$, and (ii) $\mathfrak{B} \equiv_{r, \mathcal{L}} \mathfrak{A}$. Since $\mathfrak{A} \in \mathcal{S}$ and \mathcal{S} is defined over \mathcal{S}_i by an \mathcal{L} sentence of rank r , we have $\mathfrak{B} \in \mathcal{S}$. Then \mathcal{L} -EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, k, m, \bar{a}, \theta_{(\mathcal{S}, k, \mathcal{L})}$) is true.

Suppose $m > r$. Then there exists $\mathfrak{B} \in \mathcal{S}_i$ such that \mathcal{L} -EBSP-condition($\mathcal{S}_i, \mathfrak{A}, \mathfrak{B}, k, m, \bar{a}, \theta_{(\mathcal{S}_i, k, \mathcal{L})}$) is true. Then since $\text{tp}_{\mathfrak{B}, \bar{a}, m, \mathcal{L}}(\bar{x}) = \text{tp}_{\mathfrak{A}, \bar{a}, m, \mathcal{L}}(\bar{x})$ and $m > r$, we have $\mathfrak{B} \equiv_{r, \mathcal{L}} \mathfrak{A}$ whereby reasoning as before, we have $\mathfrak{B} \in \mathcal{S}$. Then \mathcal{L} -EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, k, m, \bar{a}, \theta_{(\mathcal{S}, k, \mathcal{L})}$) is true.

Finally, since $\bar{\mathcal{S}}$ is also definable over \mathcal{S}_i by an \mathcal{L} sentence of rank r , namely by the negation of the sentence defining \mathcal{S} over \mathcal{S}_i , we have that \mathcal{L} -EBSP($\bar{\mathcal{S}}, k$) is true, with the witness function $\theta_{(\bar{\mathcal{S}}, k, \mathcal{L})}$ as in the statement of this part. \square

10.4.2 Closure of \mathcal{L} -EBSP(\cdot, \cdot) under operations implemented using translation schemes

We look at ways of generating new classes of structures that satisfy \mathcal{L} -EBSP(\cdot, \cdot) from those known to satisfy this property. Examples of well-known operations that produce new structures from given ones include “sum-like” operations [57] like disjoint union and join [19] and “product-like” operations like the cartesian and tensor products. All of these are examples of operations that are “implementable” using quantifier-free translation schemes. Let us look at the cartesian product as an example. For a vocabulary τ , let $\tau_{\text{disj-un}, 2}$ be the vocabulary ob-

tained by expanding τ with 2 fresh unary predicates P_1 and P_2 . Given structures \mathfrak{A}_1 and \mathfrak{A}_2 whose cartesian product we intend to take, we first construct the *2-disjoint sum* [38] of \mathfrak{A}_1 and \mathfrak{A}_2 , denoted $\mathfrak{A}_1 \oplus \mathfrak{A}_2$, which is the $\tau_{\text{disj-un},2}$ -structure obtained upto isomorphism, by expanding the disjoint union $\mathfrak{A}_1 \sqcup \mathfrak{A}_2$ with P_1 and P_2 interpreted respectively as the universes of the isomorphic copies of \mathfrak{A}_1 and \mathfrak{A}_2 that are used in constructing $\mathfrak{A}_1 \sqcup \mathfrak{A}_2$. The cartesian product $\mathfrak{A}_1 \otimes \mathfrak{A}_2$ is then the structure $\Xi(\mathfrak{A}_1 \oplus \mathfrak{A}_2)$ where Ξ is the $(2, \tau_{\text{disj-un},2}, \tau, \text{FO})$ -translation scheme given by $\Xi = (\xi, (\xi_R)_{R \in \tau})$ where $\xi(x, y) = (P_1(x) \wedge P_2(y))$ and for $R \in \tau$ of arity r , we have $\xi_R(x_1, y_1, \dots, x_r, y_r) = R(x_1, \dots, x_r) \wedge R(y_1, \dots, y_r)$. As a second example, consider the *across-connect* operation which takes two copies of a graph G and connects corresponding nodes across. To implement this operation, we first construct the *2-copy* of G [9]. Specifically, we take isomorphic copies G_1 and G_2 of G , where the universe of G_i is $\{(i, a) \mid a \in U_G\}$ for $i \in \{1, 2\}$. We then expand $G_1 \oplus G_2$ with the relation \sim interpreted as the set $\{((1, a), (2, a)) \mid a \in U_G\}$, to get a $\tau_{\text{copy},2}$ -structure $2\text{-copy}(G)$, where $\tau_{\text{copy},2} = \tau_{\text{disj-un},2} \cup \{\sim\}$. Then the *across-connect* of G is the structure $\Phi(2\text{-copy}(G))$ where $\Phi = (\phi, \phi_E)$, ϕ is the formula $(x = x)$ and $\phi_E(x, y) = E(x, y) \vee (P_1(x) \wedge P_2(y) \wedge (x \sim y))$. Observe that both of the translation schemes above are quantifier-free. The above operations are two instances of several useful and well-studied operations on structures that are implementable using quantifier-free translation schemes. We consider such operations in this section. Specifically, the quantifier-free translation schemes implementing the operations are those that “act” on the *n-disjoint sums* or *n-copies* of structures of a given class. We define these notions formally below.

Definition 10.4.2 (*n-disjoint sum*). Given a vocabulary τ , let τ_{k_i} be the vocabulary obtained by expanding τ with k_i fresh constant symbols, for $k_i \geq 0$ and $i \in \{1, \dots, n\}$. Let $\tau_{\text{disj-un},k_1,\dots,k_n}$ be the vocabulary obtained by expanding τ with $k_1 + \dots + k_n$ fresh constant symbols, and n fresh unary relation symbols P_1, \dots, P_n . For $i \in \{1, \dots, n\}$, let $(\mathfrak{A}_i, \bar{a}_i)$ be a τ_{k_i} -structure, where \mathfrak{A}_i is a τ -structure. Then the *n-disjoint sum* of $(\mathfrak{A}_1, \bar{a}_1), \dots, (\mathfrak{A}_n, \bar{a}_n)$, denoted $\bigoplus_{i=1}^n (\mathfrak{A}_i, \bar{a}_i)$, is the $\tau_{\text{disj-un},k_1,\dots,k_n}$ -structure \mathfrak{A} defined as follows.

1. If $\mathfrak{A}_1, \dots, \mathfrak{A}_n$ have disjoint universes, then \mathfrak{A} is such that (i) the τ -reduct of \mathfrak{A} is the disjoint union of $\mathfrak{A}_1, \dots, \mathfrak{A}_n$, (ii) $P_i^{\mathfrak{A}} = U_{\mathfrak{A}_i}$ for each $i \in \{1, \dots, n\}$ (thus the $P_i^{\mathfrak{A}}$ s form a partition of the universe of \mathfrak{A}), and (iii) for $i \in \{1, \dots, n\}$, if $l_i = k_1 + \dots + k_{i-1}$ and $l_1 = 0$, then $(c_{l_i+1}^{\mathfrak{A}}, \dots, c_{l_i+k_i}^{\mathfrak{A}}) = \bar{a}_i$, where $c_1, \dots, c_{k_1+\dots+k_n}$ are the fresh constant symbols of $\tau_{\text{disj-un},k_1,\dots,k_n}$.
2. In case, $\mathfrak{A}_1, \dots, \mathfrak{A}_n$ do not have disjoint universes, then let $(\mathfrak{A}'_i, \bar{a}'_i)$ be an isomorphic copy

of $(\mathfrak{A}_i, \bar{a}_i)$ for $i \in \{1, \dots, n\}$, such that $\mathfrak{A}'_1, \dots, \mathfrak{A}'_n$ have disjoint universes. Then \mathfrak{A} is defined upto isomorphism as the τ -structure $\bigoplus_{i=1}^{i=n} (\mathfrak{A}'_i, \bar{a}'_i)$.

Definition 10.4.3 (*n-copy*). Given a vocabulary τ , let $\tau_{\text{copy}, k_1, \dots, k_n} = \tau_{\text{disj-un}, k_1, \dots, k_n} \cup \{\sim\}$, where \sim is a binary relation symbol not in τ , and $k_1, \dots, k_n \geq 0$. Given a τ -structure \mathfrak{A} and a k_i -tuple \bar{a}_i from \mathfrak{A} for each $i \in \{1, \dots, n\}$, let $(\mathfrak{A}_i, \bar{b}_i)$ be an isomorphic copy of $(\mathfrak{A}, \bar{a}_i)$, with universe $\{(i, a) \mid a \in U_{\mathfrak{A}}\}$. Then the *n-copy of \mathfrak{A} with $\bar{a}_1, \dots, \bar{a}_n$* , denoted $n\text{-copy}(\mathfrak{A}, \bar{a}_1, \dots, \bar{a}_n)$, is the $\tau_{\text{copy}, k_1, \dots, k_n}$ -structure defined as below:

1. If $n = 1$, then $n\text{-copy}(\mathfrak{A}, \bar{a}_1) = (\mathfrak{A}, \bar{a}_1)$.
2. If $n > 1$, then $n\text{-copy}(\mathfrak{A}, \bar{a}_1, \dots, \bar{a}_n)$ is such that (i) the $\tau_{\text{disj-un}, k_1, \dots, k_n}$ -reduct of $n\text{-copy}(\mathfrak{A}, \bar{a}_1, \dots, \bar{a}_n)$ is the structure $\bigoplus_{i=1}^{i=n} (\mathfrak{A}_i, \bar{a}_i)$, and (ii) \sim is interpreted in $n\text{-copy}(\mathfrak{A}, \bar{a}_1, \dots, \bar{a}_n)$ as the set $\{((i, a), (j, a)) \mid 1 \leq i, j \leq n, a \in U_{\mathfrak{A}}\}$.

The above definitions, which are given for structures expanded with tuples of elements, instead of simply for structures that are not expanded with tuples of elements, are given so because in the proofs of our results below, we will need these general definitions. However, for the statements of our results, we deal with *n-disjoint sums* and *n-copies* of only structures that are not expanded with tuples of elements, i.e. for the case when $k_1 = \dots = k_n = 0$ in the definitions above. In such a case, we denote $\tau_{\text{disj-un}, k_1, \dots, k_n}$ and $\tau_{\text{copy}, k_1, \dots, k_n}$, simply as $\tau_{\text{disj-un}, n}$ and $\tau_{\text{copy}, n}$ respectively. Likewise, we denote $\bigoplus_{i=1}^{i=n} (\mathfrak{A}_i, \bar{a}_i)$ simply as $\bigoplus_{i=1}^{i=n} \mathfrak{A}_i$ (since each \bar{a}_i is empty), and $n\text{-copy}(\mathfrak{A}, \bar{a}_1, \dots, \bar{a}_n)$ simply as $n\text{-copy}(\mathfrak{A})$.

Given classes $\mathcal{S}_1, \dots, \mathcal{S}_n$ of τ -structures, let $n\text{-disjoint-sum}(\mathcal{S}_1, \dots, \mathcal{S}_n) = \{\bigoplus_{i=1}^{i=n} \mathfrak{A}_i \mid \mathfrak{A}_i \in \mathcal{S}_i, 1 \leq i \leq n\}$. Given a quantifier-free $(t, \tau_{\text{disj-un}, n}, \tau, \text{FO})$ -translation scheme Ξ_1 , let $\Xi_1(n\text{-disjoint-sum}(\mathcal{S}_1, \dots, \mathcal{S}_n)) = \{\Xi_1(\bigoplus_{i=1}^{i=n} \mathfrak{A}_i) \mid \mathfrak{A}_i \in \mathcal{S}_i, 1 \leq i \leq n\}$. Then Ξ_1 gives rise to an n -ary operation $O_1 : \mathcal{S}_1 \times \dots \times \mathcal{S}_n \rightarrow \Xi_1(n\text{-disjoint-sum}(\mathcal{S}_1, \dots, \mathcal{S}_n))$ defined as $O_1(\mathfrak{A}_1, \dots, \mathfrak{A}_n) = \Xi_1(\bigoplus_{i=1}^{i=n} \mathfrak{A}_i)$. Likewise, given a class \mathcal{S} of structures, if $n\text{-copy}(\mathcal{S}) = \{n\text{-copy}(\mathfrak{A}) \mid \mathfrak{A} \in \mathcal{S}\}$, then a quantifier-free $(t, \tau_{\text{copy}, n}, \tau, \text{FO})$ -translation scheme Ξ_2 gives rise to a unary operation $O_2 : \mathcal{S} \rightarrow \Xi_2(n\text{-copy}(\mathcal{S}))$ where $\Xi_2(n\text{-copy}(\mathcal{S})) = \{\Xi_2(n\text{-copy}(\mathfrak{A})) \mid \mathfrak{A} \in \mathcal{S}\}$ such that $O_2(\mathfrak{A}) = \Xi_2(n\text{-copy}(\mathfrak{A}))$. For the above cases, we say that O_1 , resp. O_2 , is *implementable using Ξ_1 , resp. Ξ_2* . We say an operation is implementable using a quantifier-free translation scheme if it is one of the two kinds of operations O_1 and O_2 just described. The following two results, which are our central results of this section, together show that operations that are implementable using quantifier-free translation schemes preserve the \mathcal{L} -EBSP(\cdot, \cdot) property of the classes they operate on.

Lemma 10.4.4. *Let $\mathcal{S}, \mathcal{S}_1, \dots, \mathcal{S}_n$ be classes of structures for $n \geq 1$. The following are true.*

1. *If \mathcal{L} -EBSP(\mathcal{S}_i, k_i) is true for $k_i \in \mathbb{N}$ for each $i \in \{1, \dots, n\}$, then so is \mathcal{L} -EBSP(n -disjoint-sum($\mathcal{S}_1, \dots, \mathcal{S}_n$), l), where $l = \min\{k_i \mid i \in \{1, \dots, n\}\}$. Further, if there is a computable witness function for \mathcal{L} -EBSP(\mathcal{S}_i, k_i) for each $i \in \{1, \dots, n\}$, then there is a computable witness function for \mathcal{L} -EBSP(n -disjoint-sum($\mathcal{S}_1, \dots, \mathcal{S}_n$), l) as well.*
2. *If \mathcal{L} -EBSP(\mathcal{S}, k) is true for $k \in \mathbb{N}$, then so is \mathcal{L} -EBSP(n -copy(\mathcal{S}), k). Further, if there is a computable witness function for \mathcal{L} -EBSP(\mathcal{S}, k), then there is a computable witness function for \mathcal{L} -EBSP(n -copy(\mathcal{S}), k) as well.*

Theorem 10.4.5. *Let \mathcal{S} be class of τ -structures, and let $\Xi = (\xi, (\xi_R)_{R \in \sigma})$ be a quantifier-free (t, τ, σ, FO) -translation scheme. Then the following hold for each $k \in \mathbb{N}$.*

1. *If FO -EBSP($\mathcal{S}, k \cdot t$) is true, then so is FO -EBSP($\Xi(\mathcal{S}), k$).*
2. *If Ξ is scalar and MSO -EBSP(\mathcal{S}, k) is true, then so is MSO -EBSP($\Xi(\mathcal{S}), k$).*

In each of the implications above, a computable witness function for the antecedent implies a computable witness function for the consequent.

Remark 10.4.6. The quantifier-freeness of Ξ in Theorem 10.4.5 is necessary in general. In fact, the presence of even a single quantifier in any one of the formulas ξ_R above can cause Theorem 10.4.5 to fail. We show this towards the end of this section.

For an operation O that is implementable using a quantifier-free translation scheme, define the *dimension* of O to be the minimum of the dimensions of the quantifier-free translation schemes that implement O . We say O is “sum-like” if its dimension is one, else we say O is “product-like”. Examples of sum-like operations include unary graph operations like complement, transpose, across-connect and the line-graph operation [19], and binary operations like disjoint union and join. Examples of product-like operations include various kinds of products such as cartesian, tensor, lexicographic, and strong products. We now have the following corollary which shows that \mathcal{L} -EBSP(\cdot, \cdot) and FO -EBSP(\cdot, \cdot) are indeed preserved under sum-like and product-like operations respectively.

Corollary 10.4.7. *Let $\mathcal{S}_1, \dots, \mathcal{S}_n, \mathcal{S}$ be classes of structures and let $O : \mathcal{S}_1 \times \dots \times \mathcal{S}_n \rightarrow \mathcal{S}$ be an n -ary operation that is implementable using a quantifier-free translation scheme. Let $O(\mathcal{S}_1, \dots, \mathcal{S}_n)$ denote the class of structures that are in the range of O , and let t be the dimension of O . Then the following are true.*

1. If \mathcal{L} -EBSP(\mathcal{S}_i, k_i) is true for $k_i \in \mathbb{N}$ for each $i \in \{1, \dots, n\}$, then so is \mathcal{L} -EBSP($\mathcal{O}(\mathcal{S}_1, \dots, \mathcal{S}_n), l$), for $l = \min\{k_i \mid i \in \{1, \dots, n\}\}$, whenever \mathcal{O} is sum-like.
2. If FO-EBSP($\mathcal{S}_i, k_i \cdot t$) is true for $k_i \in \mathbb{N}$ for each $i \in \{1, \dots, n\}$, then so is FO-EBSP($\mathcal{O}(\mathcal{S}_1, \dots, \mathcal{S}_n), l$), for $l = \min\{k_i \mid i \in \{1, \dots, n\}\}$, whenever \mathcal{O} is product-like.

In each of the implications above, if there are computable witness functions for each of the conjuncts in the antecedent, then there is a computable witness function for the consequent as well.

Proof. Follows easily from Lemma 10.4.4 and Theorem 10.4.5. \square

The rest of this section is entirely devoted to proving Lemma 10.4.4 and Theorem 10.4.5.

Towards the proof of Lemma 10.4.4, we present the following simple facts about n -disjoint sum and n -copy. We skip the proof.

Lemma 10.4.8. *Let $(\mathfrak{A}_i, \bar{a}_i)$ and $(\mathfrak{B}_i, \bar{b}_i)$ be τ_{k_i} -structures for $i \in \{1, \dots, n\}$. Let $m \in \mathbb{N}$. Then the following are true.*

1. If $(\mathfrak{B}_i, \bar{b}_i) \hookrightarrow (\mathfrak{A}_i, \bar{a}_i)$ for $i \in \{1, \dots, n\}$, then $\bigoplus_{i=1}^{i=n} (\mathfrak{B}_i, \bar{b}_i) \hookrightarrow \bigoplus_{i=1}^{i=n} (\mathfrak{A}_i, \bar{a}_i)$.
2. If $(\mathfrak{B}_i, \bar{b}_i) \equiv_{m, \mathcal{L}} (\mathfrak{A}_i, \bar{a}_i)$ for $i \in \{1, \dots, n\}$, then $\bigoplus_{i=1}^{i=n} (\mathfrak{B}_i, \bar{b}_i) \equiv_{m, \mathcal{L}} \bigoplus_{i=1}^{i=n} (\mathfrak{A}_i, \bar{a}_i)$.

Lemma 10.4.9. *Let $(\mathfrak{A}, \bar{a}_i)$ and $(\mathfrak{B}, \bar{b}_i)$ be τ_{k_i} -structures for $i \in \{1, \dots, n\}$. Let $m \in \mathbb{N}$. If $\mathcal{C} = n\text{-copy}(\mathfrak{B}, \bar{b}_1, \dots, \bar{b}_n)$ and $\mathcal{D} = n\text{-copy}(\mathfrak{A}, \bar{a}_1, \dots, \bar{a}_n)$, then the following are true.*

1. If $(\mathfrak{B}, \bar{b}_i) \hookrightarrow (\mathfrak{A}, \bar{a}_i)$ for $i \in \{1, \dots, n\}$, then $\mathcal{C} \hookrightarrow \mathcal{D}$.
2. If $(\mathfrak{B}, \bar{b}_i) \equiv_{m, \mathcal{L}} (\mathfrak{A}, \bar{a}_i)$ for $i \in \{1, \dots, n\}$, then $\mathcal{C} \equiv_{m, \mathcal{L}} \mathcal{D}$.

We now prove Lemma 10.4.4.

Proof of Lemma 10.4.4. Part 1: Consider a structure $\mathfrak{A} = (\bigoplus_{i=1}^{i=n} \mathfrak{A}_i) \in n\text{-disjoint-sum}(\mathcal{S}_1, \dots, \mathcal{S}_n)$ and let \bar{a} be an l -tuple from \mathfrak{A} . Let \bar{a}_i be the sub-tuple of \bar{a} consisting of all elements of \bar{a} that belong to $\mathcal{U}_{\mathfrak{A}_i}$; clearly $|\bar{a}_i| \leq k_i$ for $i \in \{1, \dots, n\}$. Let $m \in \mathbb{N}$. Since \mathcal{L} -EBSP(\mathcal{S}_i, k_i) is true, there exists \mathfrak{B}_i such that \mathcal{L} -EBSP-condition($\mathcal{S}_i, \mathfrak{A}_i, \mathfrak{B}_i, k_i, m, \bar{a}_i, \theta_{(\mathcal{S}_i, m, \mathcal{L})}$) holds where $\theta_{(\mathcal{S}_i, m, \mathcal{L})}$ is a witness function for \mathcal{L} -EBSP(\mathcal{S}_i, k_i). Then $(\mathfrak{B}_i, \bar{a}_i) \subseteq (\mathfrak{A}_i, \bar{a}_i)$ and $(\mathfrak{B}_i, \bar{a}_i) \equiv_{m, \mathcal{L}} (\mathfrak{A}_i, \bar{a}_i)$. Then by Lemma 10.4.8, we have that (i) $\bigoplus_{i=1}^{i=n} (\mathfrak{B}_i, \bar{b}_i) \hookrightarrow \bigoplus_{i=1}^{i=n} (\mathfrak{A}_i, \bar{a}_i)$, and (ii) $\bigoplus_{i=1}^{i=n} (\mathfrak{B}_i, \bar{b}_i) \equiv_{m, \mathcal{L}} \bigoplus_{i=1}^{i=n} (\mathfrak{A}_i, \bar{a}_i)$. Then it is easy to verify that (i) $((\bigoplus_{i=1}^{i=n} \mathfrak{B}_i), \bar{a}) \hookrightarrow ((\bigoplus_{i=1}^{i=n} \mathfrak{A}_i), \bar{a})$, and (ii) $((\bigoplus_{i=1}^{i=n} \mathfrak{B}_i), \bar{a}) \equiv_{m, \mathcal{L}} ((\bigoplus_{i=1}^{i=n} \mathfrak{A}_i), \bar{a})$. Observe that $(\bigoplus_{i=1}^{i=n} \mathfrak{B}_i) \in n\text{-disjoint-sum}(\mathcal{S}_1, \dots, \mathcal{S}_n)$, and that $|(\bigoplus_{i=1}^{i=n} \mathfrak{B}_i)| \leq \theta(m) = \sum_{i=0}^{i=n} \theta_{(\mathcal{S}_i, m, \mathcal{L})}(m)$. Taking (\mathfrak{B}, \bar{a}) to be the substructure of

(\mathfrak{A}, \bar{a}) that is isomorphic to $((\bigoplus_{i=1}^{i=n} \mathfrak{B}_i), \bar{a})$, we see \mathcal{L} -EBSP-condition(n -disjoint-sum($\mathcal{S}_1, \dots, \mathcal{S}_n$), $\mathfrak{A}, \mathfrak{B}, l, m, \bar{a}, \theta$) is true with witness function θ . Whereby \mathcal{L} -EBSP(n -disjoint-sum($\mathcal{S}_1, \dots, \mathcal{S}_n$), l) is true. It is easy to see that if $\theta_{(\mathcal{S}_i, m, \mathcal{L})}$ is computable for each $i \in \{1, \dots, n\}$, then so is θ .

Part 2: This is proved analogously as the previous part, and using Lemma 10.4.9. \square

We now proceed to proving Theorem 10.4.5. Towards the proof, we first prove the following result that shows that quantifier-free translation schemes preserve the substructure relation between any two structures of \mathcal{S} . We use results mentioned in Section 7.4 in our proof.

Lemma 10.4.10. *Let \mathcal{S} be a given class of finite structures. Let $\Xi = (\xi, (\xi_R)_{R \in \sigma})$ be a quantifier-free (t, τ, σ, FO) -translation scheme. Let \mathfrak{A} and \mathfrak{B} be given structures from \mathcal{S} , and let $\bar{b}_1, \dots, \bar{b}_n$ be n elements from $\Xi(\mathfrak{B})$, for some $n \geq 0$. If $(\mathfrak{B}, \bar{b}_1, \dots, \bar{b}_n) \subseteq (\mathfrak{A}, \bar{b}_1, \dots, \bar{b}_n)$, then (i) $\bar{b}_1, \dots, \bar{b}_n$ belong to $\Xi(\mathfrak{A})$ and (ii) $(\Xi(\mathfrak{B}), \bar{b}_1, \dots, \bar{b}_n) \subseteq (\Xi(\mathfrak{A}), \bar{b}_1, \dots, \bar{b}_n)$.*

Proof. Consider any element of $\Xi(\mathfrak{B})$; it is a t -tuple \bar{b} of \mathfrak{B} such that $(\mathfrak{B}, \bar{b}) \models \xi(\bar{x})$. Since $\xi(\bar{x})$ is quantifier-free, it is preserved under extensions over \mathcal{S} . Whereby $(\mathfrak{A}, \bar{b}) \models \xi(\bar{x})$; then \bar{b} is an element of $\Xi(\mathfrak{A})$. Since \bar{b} is an arbitrary element of $\Xi(\mathfrak{B})$, we have $\bigcup_{\Xi(\mathfrak{B})} \subseteq \bigcup_{\Xi(\mathfrak{A})}$. In particular therefore, $\bar{b}_1, \dots, \bar{b}_n$ belongs to $\Xi(\mathfrak{A})$.

Consider a relation symbol $R \in \sigma$ of arity say n . Let $\bar{d}_1, \dots, \bar{d}_n$ be elements of $\Xi(\mathfrak{B})$. Then we have the following. Below $\bar{x}_i = (x_{i,1}, \dots, x_{i,t})$ for each $i \in \{1, \dots, n\}$.

$$\begin{aligned} (\Xi(\mathfrak{B}), \bar{d}_1, \dots, \bar{d}_n) &\models R(x_1, \dots, x_n) \\ \text{iff } (\mathfrak{B}, \bar{d}_1, \dots, \bar{d}_n) &\models \Xi(R)(\bar{x}_1, \dots, \bar{x}_n) && \text{(by Proposition 7.4.2)} \\ \text{iff } (\mathfrak{B}, \bar{d}_1, \dots, \bar{d}_n) &\models \bigwedge_{i=1}^{i=n} \xi(\bar{x}_i) \wedge \xi_R(\bar{x}_1, \dots, \bar{x}_n) && \text{(by defn. of } \Xi(R); \text{ see Section 7.4)} \end{aligned}$$

Now since (i) each of ξ and ξ_R is quantifier-free, (ii) a finite conjunction of quantifier-free formulae is a quantifier-free formula, and (iii) a quantifier-free formula is preserved under substructures as well as preserved under extensions over any class, we have that

$$\begin{aligned} (\mathfrak{B}, \bar{d}_1, \dots, \bar{d}_n) &\models \bigwedge_{i=1}^{i=n} \xi(\bar{x}_i) \wedge \xi_R(\bar{x}_1, \dots, \bar{x}_n) \\ \text{iff } (\mathfrak{A}, \bar{d}_1, \dots, \bar{d}_n) &\models \bigwedge_{i=1}^{i=n} \xi(\bar{x}_i) \wedge \xi_R(\bar{x}_1, \dots, \bar{x}_n) \\ \text{iff } (\mathfrak{A}, \bar{d}_1, \dots, \bar{d}_n) &\models \Xi(R)(\bar{x}_1, \dots, \bar{x}_n) && \text{(by definition of } \Xi(R)) \\ \text{iff } (\Xi(\mathfrak{A}), \bar{d}_1, \dots, \bar{d}_n) &\models R(x_1, \dots, x_n) && \text{(by Proposition 7.4.2)} \end{aligned}$$

Since R is an arbitrary relation symbol of σ , we have that $\Xi(\mathfrak{B}) \subseteq \Xi(\mathfrak{A})$, whereby $(\Xi(\mathfrak{B}), \bar{d}_1, \dots, \bar{d}_n) \subseteq (\Xi(\mathfrak{A}), \bar{d}_1, \dots, \bar{d}_n)$. \square

Proof of Theorem 10.4.5. Part 1: Consider a structure $\Xi(\mathfrak{A}) \in \Xi(\mathcal{S})$ for some structure $\mathfrak{A} \in \mathcal{S}$. Let $(\bar{a}_1, \dots, \bar{a}_k)$ be a k -tuple from $\Xi(\mathfrak{A})$ and let $m \in \mathbb{N}$. For each $i \in \{1, \dots, k\}$, let $\bar{a}_i = (a_{i,1}, \dots, a_{i,t})$. Let $p = k \cdot t$ and consider the p -tuple \bar{a} from \mathfrak{A} given by $\bar{a} = (a_{1,1}, \dots, a_{1,t}, a_{2,1}, \dots, a_{2,t}, \dots, a_{k,1}, \dots, a_{k,t})$. Let $r = t \cdot m$. Since $\text{FO-EBSP}(\mathcal{S}, p)$ is true, there exists a witness function $\theta_{(\mathcal{S}, p, \text{FO})} : \mathbb{N} \rightarrow \mathbb{N}$ and a structure \mathfrak{B} such that $\text{FO-EBSP-condition}(\mathcal{S}, \mathfrak{A}, \mathfrak{B}, p, r, \bar{a}, \theta_{(\mathcal{S}, p, \text{FO})})$ is true. That is (i) $\mathfrak{B} \in \mathcal{S}$, (ii) $\mathfrak{B} \subseteq \mathfrak{A}$, (iii) the elements of \bar{a} are contained in \mathfrak{B} , (iv) $|\mathfrak{B}| \leq \theta_{(\mathcal{S}, p, \text{FO})}(r)$ and (v) $(\mathfrak{B}, \bar{a}) \equiv_r (\mathfrak{A}, \bar{a})$.

We now show that there exists a function $\theta_{(\Xi(\mathcal{S}), k, \text{FO})} : \mathbb{N} \rightarrow \mathbb{N}$ such that $\text{FO-EBSP-condition}(\Xi(\mathcal{S}), \Xi(\mathfrak{A}), \Xi(\mathfrak{B}), k, m, (\bar{a}_1, \dots, \bar{a}_k), \theta_{(\Xi(\mathcal{S}), k, \text{FO})})$ is true. This would show that $\text{FO-EBSP}(\Xi(\mathcal{S}), k)$ is true.

- (i) $\Xi(\mathfrak{B}) \in \Xi(\mathcal{S})$: Obvious from the definition of $\Xi(\mathcal{S})$ and the fact that $\mathfrak{B} \in \mathcal{S}$.
- (ii) $\Xi(\mathfrak{B}) \subseteq \Xi(\mathfrak{A})$: Follows from Lemma 10.4.10.
- (iii) The element \bar{a}_i is contained in $\text{U}_{\Xi(\mathfrak{B})}$ for each $i \in \{1, \dots, k\}$: Since the elements of \bar{a} are contained in \mathfrak{B} , we have for each $i \in \{1, \dots, k\}$, that \bar{a}_i is a t -tuple from \mathfrak{B} . Now since \bar{a}_i is an element of $\Xi(\mathfrak{A})$, we have $(\mathfrak{A}, \bar{a}_i) \models \xi(\bar{x})$. Since $\xi(\bar{x})$ is quantifier-free, it is preserved under substructures over \mathcal{S} . Whereby $(\mathfrak{B}, \bar{a}_i) \models \xi(\bar{x})$; then \bar{a}_i is an element of $\Xi(\mathfrak{B})$, for each $i \in \{1, \dots, k\}$.
- (iv) $(\Xi(\mathfrak{B}), \bar{a}_1, \dots, \bar{a}_k) \equiv_m (\Xi(\mathfrak{A}), \bar{a}_1, \dots, \bar{a}_k)$: Since $(\mathfrak{B}, \bar{a}) \equiv_r (\mathfrak{A}, \bar{a})$, it follows from Corollary 7.4.3, that $(\Xi(\mathfrak{B}), \bar{a}_1, \dots, \bar{a}_k) \equiv_m (\Xi(\mathfrak{A}), \bar{a}_1, \dots, \bar{a}_k)$.
- (v) The existence of a function $\theta_{(\Xi(\mathcal{S}), k, \text{FO})} : \mathbb{N} \rightarrow \mathbb{N}$ such that $|\Xi(\mathfrak{B})| \leq \theta_{(\Xi(\mathcal{S}), k, \text{FO})}(m)$: Define $\theta_{(\Xi(\mathcal{S}), k, \text{FO})} : \mathbb{N} \rightarrow \mathbb{N}$ as $\theta_{(\Xi(\mathcal{S}), k, \text{FO})}(m) = (\theta_{(\mathcal{S}, p, \text{FO})}(r))^t$. Since $|\mathfrak{B}| \leq \theta_{(\mathcal{S}, p, \text{FO})}(r)$, we have that $|\Xi(\mathfrak{B})| \leq \theta_{(\Xi(\mathcal{S}), k, \text{FO})}(m)$.

It is clear that if $\theta_{(\mathcal{S}, p, \text{FO})}$ is computable, then so is $\theta_{(\Xi(\mathcal{S}), k, \text{FO})}$.

Part 2: The proof of this part is similar to the proof above. □

Necessity of the condition on Ξ of being quantifier-free in Theorem 10.4.5:

Let $\tau = \{\leq\}$ and $\sigma = \{E\}$ where \leq, E are binary relation symbols. Consider the $(1, \tau, \sigma, \text{FO})$ -translation scheme Ξ_1 given by $\Xi_1 = (\xi_1, \xi_E^1)$ where $\xi_1(x)$ is the formula $(x = x)$ and $\xi_E^1(x, y) = \forall z((x \leq z) \wedge (x \neq z)) \rightarrow (y \leq z)$. Consider the class \mathcal{S} of all finite linear orders. We know from Theorem 10.2.2 that both $\text{FO-EBSP}(\mathcal{S}, l)$ and $\text{MSO-EBSP}(\mathcal{S}, l)$ hold for all $l \in \mathbb{N}$. The (universal) formula $\xi_E^1(x, y)$ cannot be \mathcal{S} -equivalent to a quantifier-free formula. To see this, suppose $\xi_E^1(x, y)$ is \mathcal{S} -equivalent to a quantifier-free formula $\beta(x, y)$. Consider

the structure $\mathfrak{A} = (\{1, 2, 3\}, \leq^{\mathfrak{A}}) \in \mathcal{S}$ where $\leq^{\mathfrak{A}}$ is the usual linear order on $\{1, 2, 3\}$. Clearly $(\mathfrak{A}, 1, 3) \models \neg \xi_E^1(x, y)$ whereby $(\mathfrak{A}, 1, 3) \models \neg \beta(x, y)$. Since $\neg \beta$ is quantifier-free, it is preserved under substructures whereby $\mathfrak{B} = (\{1, 3\}, \{(1, 1), (1, 3), (3, 3)\})$ is such that $\mathfrak{B} \models \neg \beta(x, y)$ and hence $\mathfrak{B} \models \neg \xi_E^1(x, y)$. The latter is clearly not true. Then Ξ_1 is not quantifier-free.

We now show that $\text{FO-EBSP}(\Xi_1(\mathcal{S}), k)$ is false for each $k \geq 2$. The class $\Xi_1(\mathcal{S})$ is the class of all finite directed paths. It is easy to see that for $m, n \in \mathbb{N}$ such that $m \geq 4$ and $n \geq 2$, the path P_n of length n (i.e. having $n + 1$ vertices) is not m -equivalent to any substructure of P_n that contains both the end-points of P_n and that has size at most n . Then $\text{FO-EBSP}(\Xi_1(\mathcal{S}), k)$ is false for each $k \geq 2$.

Consider the $(1, \tau, \sigma, \text{FO})$ -translation scheme Ξ_2 given by $\Xi_2 = (\xi_2, \xi_E^2)$ where $\xi_2 = \xi_1$ and $\xi_E^2 = \neg \xi_E^1$. Let $\text{Neg}_\sigma = (\alpha, \alpha_E)$ be the $(1, \sigma, \sigma, \text{FO})$ -translation scheme that is quantifier-free and such that $\alpha(x)$ is the formula $(x = x)$ and $\alpha_E(x, y) = \neg E(x, y)$. For the class \mathcal{S} as above, observe that $\Xi_1(\mathcal{S})$ is exactly the class $\text{Neg}_\sigma(\Xi_2(\mathcal{S}))$. Whence if $\text{FO-EBSP}(\Xi_2(\mathcal{S}), k)$ is true for some $k \geq 2$, then by Part (1) above, $\text{FO-EBSP}(\text{Neg}_\sigma(\Xi_2(\mathcal{S})), k)$ is true, contradicting the fact that $\text{FO-EBSP}(\Xi_1(\mathcal{S}), k)$ is false for all $k \geq 2$.

10.4.3 Closure under regular operation-tree languages

Theorem 10.4.5 shows us that operations that are implemented using quantifier-free translation schemes, preserve the $\text{FO-EBSP}(\cdot, \cdot)$ or $\text{MSO-EBSP}(\cdot, k)$ property of the class of structures they are applied to. From this, and from Lemma 10.4.1(2), it follows that finite unions of the classes obtained by applying finite compositions of the aforesaid kind of operations to a given class \mathcal{S} of structures, also preserves the $\text{FO-EBSP}(\cdot, \cdot)$ or $\text{MSO-EBSP}(\cdot, k)$ property of \mathcal{S} . However, as already mentioned in the introduction, there are interesting classes of structures that are produced only by taking infinite such unions; examples include hamming graphs of the n -clique, and the class of all p -dimensional grid posets, where p belongs to an MSO-definable (using a linear order) class of natural numbers. In this section, we discuss the case of such infinite unions. Specifically, we show that under reasonable additional assumptions on the aforementioned operations, that are satisfied by the operations of disjoint union, join, across connect, and the various kinds of products mentioned in Section 10.4.2, it is the case that the property of \mathcal{L} -EBSP($\cdot, 0$) of a class is preserved under taking the aforementioned infinite unions, provided that these unions are “regular” in a sense that we make precise. Indeed the infinite unions that

produce the examples of hamming graphs of the n -clique, and the class of p -dimensional grid posets referred to above, are regular in our sense, whereby since the examples are produced using the cartesian product operation, it follows that each of these satisfies \mathcal{L} -EBSP($\cdot, 0$).

Let Op be a finite set of operations implementable using quantifier-free translation schemes. We call the operations of Op as *quantifier-free operations*, and abusing notation, use Ξ to represent these operations. Let $\rho : \text{Op} \rightarrow \mathbb{N}$ be such that $\rho(\Xi)$ is the arity of Ξ , for $\Xi \in \text{Op}$. An *operation tree over* Op is an ordered tree ranked by ρ , in which each internal node is labeled with an operation of Op , and each leaf node is labeled with the label \diamond , which is a place-holder for an “input” structure. The singleton tree (without any internal nodes) in which the sole node is labeled with a \diamond is also an operation tree over Op (treated as the “no operation” tree). When the \diamond labels of the leaf nodes of an operation tree t are replaced with structures, then the resulting tree s can naturally be seen as a representation tree of a structure \mathfrak{A}_s . Formally, the structure \mathfrak{A}_s can be defined (up to isomorphism) inductively as follows. If s is a singleton, then \mathfrak{A}_s is the structure labeling the sole node of s . Else, let a_1, \dots, a_n be in increasing order, the children of the root of s . Let $t_i = s_{\geq a_i}$ be the subtree of s rooted at a_i , for $i \in \{1, \dots, n\}$. Assume (as induction hypothesis) that the structure \mathfrak{A}_{t_i} represented (upto isomorphism) by the tree t_i is already defined. Let Ξ be the operation labeling the root of s . Then $\mathfrak{A}_s = \Xi(\mathfrak{A}_{t_1}, \dots, \mathfrak{A}_{t_n})$ upto isomorphism.

Given an operation tree t over Op and a class \mathcal{S} of structures, let $t(\mathcal{S})$ be the *isomorphism-closed* class of structures represented by the representation trees obtained by simply replacing the labels of the leaf nodes of t , with structures from \mathcal{S} . By extension, given a class \mathcal{V} of operation trees over Op , let $\mathcal{V}(\mathcal{S}) = \bigcup_{t \in \mathcal{V}} t(\mathcal{S})$. The class $\mathcal{V}(\mathcal{S})$ is then isomorphism-closed as well. If \mathcal{V} is finite, then Theorem 10.4.5 and Lemma 10.4.1(2) show that \mathcal{L} -EBSP(\cdot, \cdot) property of \mathcal{S} remains preserved under \mathcal{V} , where \mathcal{V} is seen as a transformation of a class of structures. While we are yet to investigate what happens if \mathcal{V} is an arbitrary infinite class, we show below that if \mathcal{V} , seen as a language of ordered ranked trees over $\text{Op} \cup \{\diamond\}$, is regular (in the sense of regularity used in the literature for ordered ranked trees), then the truth of \mathcal{L} -EBSP($\cdot, 0$) is preserved in going from \mathcal{S} to $\mathcal{V}(\mathcal{S})$, provided that the operations in Op satisfy the additional properties of “monotonicity” and “ $\equiv_{m, \mathcal{L}}$ -preservation” that we define below.

An n -ary operation Ξ is said to be *monotone* if for all structures $\mathfrak{A}_1, \dots, \mathfrak{A}_n$, we have \mathfrak{A}_i is (isomorphically) embeddable in $\Xi(\mathfrak{A}_1, \dots, \mathfrak{A}_n)$. We say Ξ is *$\equiv_{m, \mathcal{L}}$ -preserving* if for all structures $\mathfrak{A}_1, \dots, \mathfrak{A}_n, \mathfrak{B}_1, \dots, \mathfrak{B}_n$, it is the case that if $\mathfrak{A}_i \equiv_{m, \mathcal{L}} \mathfrak{B}_i$ for all $i \in \{1, \dots, n\}$, then

$\Xi(\mathfrak{A}_1, \dots, \mathfrak{A}_n) \equiv_{m, \mathcal{L}} \Xi(\mathfrak{B}_1, \dots, \mathfrak{B}_n)$. The operations of disjoint union, join and across connect seen in Section 10.4.2 are monotone and $\equiv_{m, \text{MSO}}$ -preserving, while each of the products mentioned in Section 10.4.2, like cartesian, tensor, lexicographic and strong products, is monotone and $\equiv_{m, \text{FO}}$ -preserving. The central result of this section can now be stated as follows.

Theorem 10.4.11. *Let Op be a finite set of operations, where each operation in Op is quantifier-free, monotone and $\equiv_{m, \mathcal{L}}$ -preserving. Let \mathcal{V} be a class of operation trees over Op , that is regular. Let \mathcal{S} be a class of structures. If $\mathcal{L}\text{-EBSP}(\mathcal{S}, 0)$ is true, then so is $\mathcal{L}\text{-EBSP}(\mathcal{V}(\mathcal{S}), 0)$. Further, if $\mathcal{L}\text{-EBSP}(\mathcal{S}, 0)$ has a computable witness function, then so does $\mathcal{L}\text{-EBSP}(\mathcal{V}(\mathcal{S}), 0)$.*

Proof. We assume familiarity with the notions and results of Section 10.1 for the present proof. Let $\mathcal{S}_1 = \bigcup_{t \in \mathcal{S}_2} t(\mathcal{S})$, where \mathcal{S}_2 be the class of all operation trees over Op . Then \mathcal{S} and \mathcal{V} are resp. subclasses of \mathcal{S}_1 and \mathcal{S}_2 . Let $\Sigma_{\text{int}} = \text{Op}$ and $\Sigma_{\text{leaf}} = \{\mathfrak{A} \mid \text{U}_{\mathfrak{A}} \subseteq \mathbb{N}, \mathfrak{A} \cong \mathfrak{B}, \mathfrak{B} \in \mathcal{S}\}$. Observe that Σ_{leaf} is countable. Let $\rho : \text{Op} \rightarrow \mathbb{N}$ be such that $\rho(\Xi)$ is the arity of Ξ . Let \mathcal{T} be the class of all representation-feasible trees over $\Sigma_{\text{int}} \cup \Sigma_{\text{leaf}}$, that are ranked by ρ ; then \mathcal{T} is closed under rooted subtrees and under replacements with rooted subtrees.

We now construct two representation maps $\text{Str}_i : \mathcal{T} \rightarrow \mathcal{S}_i$ for $i \in \{1, 2\}$ such that for $s \in \mathcal{T}$, $\text{Str}_1(s)$ is the structure \mathfrak{A}_s represented by s (as defined earlier), while $\text{Str}_2(s)$ is the operation tree corresponding to s (i.e. the tree obtained by simply replacing the leaf nodes of s with \diamond).

We now observe the following.

1. The map Str_1 satisfies conditions B.1 and A.1.b of Section 10.1, for each $m \in \mathbb{N}$, because each operation in Op is assumed to be monotone and $\equiv_{m, \mathcal{L}}$ preserving. That Str_1 also satisfies A.1.a is seen by observing that each operation in Op is implementable using a quantifier-free translation scheme (see the paragraph before Lemma 10.4.4 for the precise meaning of implementability using quantifier-free translation schemes), and then using Lemmas 10.4.8, 10.4.9 and 10.4.10. Whereby, Str_1 is \mathcal{L} -height-reduction favourable.
2. The map Str_2 is easily seen to satisfy conditions A.1.a and B.1. That it also satisfies A.1.b for $\mathcal{L} = \text{MSO}$ and for all $m \geq 2$ follows from the MSO composition lemma for ordered trees (see Lemma 10.2.3). Whereby, Str_2 is MSO-height-reduction favourable.

Let \mathfrak{A} be a structure in $\mathcal{V}(\mathcal{S})$, and let $m \geq 2$. We show below the existence of a structure \mathfrak{B} such that $\mathcal{L}\text{-EBSP-condition}(\mathcal{V}(\mathcal{S}), \mathfrak{A}, \mathfrak{B}, 0, m, \text{null}, \theta_{(\mathcal{V}(\mathcal{S}), 0, \mathcal{L})})$ holds, where null is the empty tuple and $\theta_{(\mathcal{V}(\mathcal{S}), 0, \mathcal{L})}$ is a function from \mathbb{N} to \mathbb{N} such that $\theta_{(\mathcal{V}(\mathcal{S}), 0, \mathcal{L})}(p) = \theta_{(\mathcal{V}(\mathcal{S}), 0, \mathcal{L})}(2)$ for $p \leq 2$. It is obvious then that $\mathcal{L}\text{-EBSP-condition}(\mathcal{V}(\mathcal{S}), \mathfrak{A}, \mathfrak{B}, 0, p, \text{null}, \theta_{(\mathcal{V}(\mathcal{S}), 0, \mathcal{L})})$ holds for $p \leq 2$. Then $\mathcal{L}\text{-EBSP}(\mathcal{V}(\mathcal{S}), 0)$ holds with $\theta_{(\mathcal{V}(\mathcal{S}), 0, \mathcal{L})}$ being a witness function.

Since $\mathfrak{A} \in \mathcal{V}(\mathcal{S})$, there exists $t \in \mathcal{T}$ such that $\text{Str}_1(t) \cong \mathfrak{A}$ and $\text{Str}_2(t) \in \mathcal{V}$. Since \mathcal{V} is regular, it is defined by an MSO sentence φ (the sentence φ exists since regularity corresponds to MSO definability for ordered ranked trees; see Section 10.2). Then $\text{Str}_2(t) \models \varphi$. Let the rank of φ be n . Since Str_1 is \mathcal{L} -height-reduction favourable and Str_2 is MSO-height-reduction favourable, we have by Theorem 10.1.1, that there is a computable function $\eta_2 : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ and a subtree s_2 of t in \mathcal{T} , such that (i) the height of s_2 is at most $\eta_2(m, n)$, (ii) $\text{Str}_1(s_2) \hookrightarrow \text{Str}_1(t)$, (iii) $\text{Str}_1(s_2) \equiv_{m, \mathcal{L}} \text{Str}_1(t)$, and (iv) $\text{Str}_2(s_2) \equiv_{n, \text{MSO}} \text{Str}_2(t)$. Since $\text{Str}_2(t) \models \varphi$, we have $\text{Str}_2(s_2) \models \varphi$ whereby $\text{Str}_2(s_2) \in \mathcal{V}$.

Now since $\mathcal{L}\text{-EBSP}(\mathcal{S}, 0)$ is true, we have for each structure $\mathfrak{C} \in \mathcal{S}$, a structure $\mathfrak{C}' \in \mathcal{S}$ such that (i) $\mathfrak{C}' \subseteq \mathfrak{C}$ (ii) $|\mathfrak{C}'| \leq \theta_{(\mathcal{S}, 0, \mathcal{L})}(m)$ and (iii) $\mathfrak{C}' \equiv_{m, \mathcal{L}} \mathfrak{C}$, where $\theta_{(\mathcal{S}, 0, \mathcal{L})}$ is a witness function for $\mathcal{L}\text{-EBSP}(\mathcal{S}, 0)$. Let $s_1 \in \mathcal{T}$ be the tree obtained from s_2 by replacing each structure \mathfrak{C} labeling a leaf of s_2 with the structure \mathfrak{C}' described above. Since Str_1 satisfies conditions A.1.a and A.1.b for each $m \geq 2$, one can verify that (i) $\text{Str}_1(s_1) \hookrightarrow \text{Str}_1(s_2) \hookrightarrow \text{Str}_1(t) \cong \mathfrak{A}$, and (ii) $\text{Str}_1(s_1) \equiv_{m, \mathcal{L}} \text{Str}_1(s_2) \equiv_{m, \mathcal{L}} \text{Str}_1(t)$. It is clear that $\text{Str}_1(s_1) \in \mathcal{V}(\mathcal{S})$ since $\text{Str}_2(s_1) = \text{Str}_2(s_2) \in \mathcal{V}$. Let \mathfrak{B} be the substructure of \mathfrak{A} such that $\mathfrak{B} \cong \text{Str}_1(s_1)$. Then from the above discussion, we have $\mathfrak{B} \equiv_{m, \mathcal{L}} \mathfrak{A}$ and $\mathfrak{B} \in \mathcal{V}(\mathcal{S})$. We now show that \mathfrak{B} is of bounded size. Let d be the maximum arity of any operation in Op , and t be the maximum of the dimensions of the translation schemes implementing the operations in Op (see the definition of dimension in Section 7.4). Recall that (i) the height of s_1 is the same as the height of s_2 which in turn is at most $\eta_2(m, n)$, and (ii) the size of any structure labeling a leaf node of s_1 is at most $\theta_{(\mathcal{S}, 0, \mathcal{L})}(m)$. Then $|\mathfrak{B}| = |\text{Str}_1(s_1)| \leq \theta_{(\mathcal{V}(\mathcal{S}), 0, \mathcal{L})}(m) = f(0)$ where for $0 \leq j \leq \eta_2(m, n)$, $f(j)$ is as defined below.

$$f(j) = \begin{cases} \theta_{(\mathcal{S}, 0, \mathcal{L})}(m) & \text{if } j = \eta_2(m, n) \\ (d \cdot f(j+1))^t & \text{if } j < \eta_2(m, n) \end{cases}$$

It is now easy to verify that $\mathcal{L}\text{-EBSP-condition}(\mathcal{V}(\mathcal{S}), \mathfrak{A}, \mathfrak{B}, 0, m, \text{null}, \theta_{(\mathcal{V}(\mathcal{S}), 0, \mathcal{L})})$ is true, where null denotes the empty tuple. Define $\theta_{(\mathcal{V}(\mathcal{S}), 0, \mathcal{L})}(p) = \theta_{(\mathcal{V}(\mathcal{S}), 0, \mathcal{L})}(2)$ for $p \leq 2$. Then as reasoned earlier, we have that $\mathcal{L}\text{-EBSP}(\mathcal{V}(\mathcal{S}), 0)$ holds with $\theta_{(\mathcal{V}(\mathcal{S}), 0, \mathcal{L})}$ being a witness function. One can see that if $\theta_{(\mathcal{S}, 0, \mathcal{L})}$ is computable, then so is $\theta_{(\mathcal{V}(\mathcal{S}), 0, \mathcal{L})}$. \square

Using the above theorem, we show below that each of the following classes, that motivated this section, satisfies $\mathcal{L}\text{-EBSP}(\cdot, 0)$: the class of hamming graphs of the n -clique, and the class of all p -dimensional grid posets where p belongs to an MSO definable (using a linear order) class of natural numbers.

1. Let \mathcal{S} be a class consisting of only the n -clique upto isomorphism. Let $\text{Op} = \{\times\}$ where \times denotes cartesian product. Let \mathcal{V} be the class of all trees over Op ; clearly \mathcal{V} is defined by the sentence True , and is trivially regular. Observe that the class $\mathcal{V}(\mathcal{S})$ is exactly the class of all hamming graphs of the n -clique. Since \mathcal{S} is finite, $\mathcal{L}\text{-EBSP}(\mathcal{S}, 0)$, and hence $\text{FO-EBSP}(\mathcal{S}, 0)$, is true with a computable witness function (see Chapter 9). Since \times is quantifier-free, monotone and $\equiv_{m, \text{FO}}$ -preserving, we have by Theorem 10.4.11 that $\text{FO-EBSP}(\mathcal{V}(\mathcal{S}), 0)$ is true with a computable witness function.
2. Let \mathcal{S} be the class of all linear orders. Let $\text{Op} = \{\times\}$. Let \mathcal{U} be the class of all operation-trees over Op in which each internal node has exactly two children, at least one of which is a leaf. It is easy to see that any tree in \mathcal{U} has a “spine” consisting of the internal nodes of the tree. Let \mathcal{V} be any MSO definable (over the class of all trees over Op) subclass of \mathcal{U} (like for instance, the class of all trees of \mathcal{U} having a spine of even length). Then \mathcal{V} is clearly regular. Since Op is a singleton, we can identify \mathcal{V} with a set Z of natural numbers that is definable in MSO using a linear order. Then $\mathcal{V}(\mathcal{S})$ can be seen as the class of all p -dimensional grid posets where $p \in Z$. Since $\mathcal{L}\text{-EBSP}(\mathcal{S}, 0)$, and hence $\text{FO-EBSP}(\mathcal{S}, 0)$, is true with a computable witness function (by Theorem 10.2.2), it follows from Theorem 10.4.11, that $\text{FO-EBSP}(\mathcal{V}(\mathcal{S}), 0)$ is also true with a computable witness function.

One can ask what happens to Theorem 10.4.11 for $k > 0$. From the very special cases we have managed to solve so far, we believe that new techniques would be necessary, in addition to the ones currently employed in proving Theorem 10.4.11.

Chapter 11

Additional studies on \mathcal{L} -EBSP(\cdot, k)

11.1 \mathcal{L} -EBSP(\mathcal{S}, k) and the decidability of \mathcal{L} -Th(\mathcal{S})

Denote by \mathcal{L} -Th(\mathcal{S}) the \mathcal{L} -theory of \mathcal{S} , i.e. the set of all \mathcal{L} sentences that are true in all structures of \mathcal{S} . We have the following result.

Lemma 11.1.1. *Let \mathcal{S} be a class of structures such that \mathcal{L} -EBSP(\mathcal{S}, k) holds for some $k \in \mathbb{N}$. If there exists a computable witness function for \mathcal{L} -EBSP(\mathcal{S}, k), then \mathcal{L} -Th(\mathcal{S}) is decidable.*

Proof. Let φ be an \mathcal{L} sentence of rank m . Let $\psi = \neg\varphi$ be the negation of φ ; then ψ has rank m as well. Suppose ψ is satisfied in a structure $\mathfrak{A} \in \mathcal{S}$. Then since \mathcal{L} -EBSP(\mathcal{S}, k) is true, for any k -tuple \bar{a} from \mathfrak{A} , there exists a structure \mathfrak{B} such that \mathcal{L} -EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, k, m, \bar{a}, \theta_{(\mathcal{S}, k, \mathcal{L})}$) is true, where $\theta_{(\mathcal{S}, k, \mathcal{L})}$ is a witness function for \mathcal{L} -EBSP(\mathcal{S}, k). Then, (i) $\mathfrak{B} \in \mathcal{S}$, (ii) $|\mathfrak{B}| \leq \theta_{(\mathcal{S}, k, \mathcal{L})}(m)$, and (iii) $\mathfrak{B} \equiv_{m, \mathcal{L}} \mathfrak{A}$. Whereby $\mathfrak{B} \models \psi$ since $\mathfrak{A} \models \psi$ and the rank of ψ is m . Thus, if ψ is satisfiable over \mathcal{S} , it is satisfied in a structure of \mathcal{S} , of size $\leq \theta_{(\mathcal{S}, k, \mathcal{L})}(m)$. Whereby, if $\theta_{(\mathcal{S}, k, \mathcal{L})}$ is a computable function, the following algorithm \mathcal{A} decides membership in \mathcal{L} -Th(\mathcal{S}).

Algorithm \mathcal{A} :

1. Compute the rank m of the input sentence φ , and compute the number $p = \theta_{(\mathcal{S}, k, \mathcal{L})}(m)$.
2. Enumerate all the finitely many structures \mathfrak{C} in \mathcal{S} of size $\leq p$, and check if the sentence $\psi = \neg\varphi$ is true in all of them. Checking if ψ is true in \mathfrak{C} is effective since \mathfrak{C} is finite.
3. If some structure \mathfrak{C} is found satisfying ψ in the previous step, then output “No”, else output “Yes”.

It is clear that \mathcal{A} indeed decides \mathcal{L} -Th(\mathcal{S}). □

As seen in Chapter 10, a wide array of classes \mathcal{S} satisfy either FO-EBSP(\mathcal{S}, k) or MSO-EBSP(\mathcal{S}, k) with computable witness functions, whereby FO-Th(\mathcal{S}) or MSO-Th(\mathcal{S}) resp., is decidable.

11.2 \mathcal{L} -EBSP(\cdot, k) and well-quasi-ordering under embedding

A pre-order (A, \leq) is said to be a *well-quasi-order* (w.q.o.) if for every infinite sequence a_1, a_2, \dots of elements of A , there exists $i < j$ such that $a_i \leq a_j$ (see [19]). If (A, \leq) is a w.q.o., we say that “ A is a w.q.o. under \leq ”. An elementary fact is that if A is a w.q.o. under \leq , then for every infinite sequence a_1, a_2, \dots of elements of A , there exists an infinite subsequence a_{i_1}, a_{i_2}, \dots such that $i_1 < i_2 < \dots$ and $a_{i_1} \leq a_{i_2} \leq \dots$.

Given a vocabulary τ and $k \in \mathbb{N}$, let τ_k be as usual, the vocabulary obtained by expanding τ with k fresh and distinct constant symbols. Let \mathcal{S} be a class of τ -structures. Denote by \mathcal{S}^k the class of all τ_k -structures whose τ -reducts are structures in \mathcal{S} . Observe that $(\mathcal{S}^k, \hookrightarrow)$ is a pre-order. We now define the property $\text{WQO}(\mathcal{S}, k)$ via the notion of w.q.o. mentioned above.

Definition 11.2.1. We say that $\text{WQO}(\mathcal{S}, k)$ holds if $(\mathcal{S}^k, \hookrightarrow)$ is a well-quasi-order.

A simple example of a class \mathcal{S} of structures satisfying $\text{WQO}(\mathcal{S}, k)$ for every $k \in \mathbb{N}$ is a finite class of finite structures. The celebrated results such as Higman’s lemma and Kruskal’s tree theorem [19] state that $\text{WQO}(\text{Words}(\Sigma), 0)$ and $\text{WQO}(\text{Unordered-trees}(\Sigma), 0)$ respectively hold. Also, the results in [31] show that $\text{WQO}(n\text{-partite-cographs}, 0)$ holds, for every $n \in \mathbb{N}$, where n -partite-cographs is the class of all n -partite cographs.

A priori, there is no reason to expect any relation between the $\text{WQO}(\cdot, k)$ and \mathcal{L} -EBSP(\cdot, k) properties. Surprisingly, we have the following result.

Theorem 11.2.2. Let \mathcal{S} be a class of structures that is closed under isomorphisms, and let $k \in \mathbb{N}$. If $\text{WQO}(\mathcal{S}, k)$ holds, then so does \mathcal{L} -EBSP(\mathcal{S}, k).

Proof. We prove the result by contradiction. Suppose, if possible, $\text{WQO}(\mathcal{S}, k)$ holds but \mathcal{L} -EBSP(\mathcal{S}, k) fails. Then by Definition 9.1, there exists $m \in \mathbb{N}$ such that for all $p \in \mathbb{N}$, there exists a structure \mathfrak{A}_p in \mathcal{S} and a k -tuple \bar{a}_p from \mathfrak{A}_p such that for any structure $\mathfrak{B} \in \mathcal{S}$, we have $((\mathfrak{B} \subseteq \mathfrak{A}_p) \wedge (\bar{a}_p \in \mathcal{U}_{\mathfrak{B}}^k) \wedge (|\mathfrak{B}| \leq p)) \rightarrow (\mathfrak{B}, \bar{a}_p) \not\equiv_{m, \mathcal{L}} (\mathfrak{A}_p, \bar{a}_p)$.

For each $p \geq 1$, fix the structure \mathfrak{A}_p and the tuple \bar{a}_p that satisfy the above properties. Let \mathfrak{A}'_p be the structure $(\mathfrak{A}_p, \bar{a}_p) \in \mathcal{S}^k$. Consider the sequence $(\mathfrak{A}'_i)_{i \geq 1}$. Since $\text{WQO}(\mathcal{S}, k)$ holds, \mathcal{S}^k is a w.q.o. under \hookrightarrow . Therefore, there exists an infinite sequence $I = (i_1, i_2, \dots)$ of indices such that $i_1 < i_2 < \dots$ and $\mathfrak{A}'_{i_1} \hookrightarrow \mathfrak{A}'_{i_2} \hookrightarrow \dots$. Consider $\Delta_{\mathcal{L}}(m, \mathcal{S}^k)$ – the set of all equivalence classes of the $\equiv_{m, \mathcal{L}}$ relation over the structures of \mathcal{S}^k . From Proposition 7.2.1, we see that $\Delta_{\mathcal{L}}(m, \mathcal{S}^k)$ is a finite set. Therefore, there exists an infinite subsequence $J = (j_1, j_2, \dots)$ of I such that (i)

$j_1 < j_2 < \dots$ (ii) $\mathfrak{A}'_{j_1} \hookrightarrow \mathfrak{A}'_{j_2} \hookrightarrow \dots$, and (iii) $\mathfrak{A}'_{j_1}, \mathfrak{A}'_{j_2}, \dots$ are all in the same $\equiv_{m, \mathcal{L}}$ class. Let $r = |\mathfrak{A}'_{j_1}|$, and let $n > 1$ be an index such that $j_n \geq r$. Then $\mathfrak{A}'_{j_1} \hookrightarrow \mathfrak{A}'_{j_n}$ and $\mathfrak{A}'_{j_1} \equiv_{m, \mathcal{L}} \mathfrak{A}'_{j_n}$. Fix an embedding $\iota : \mathfrak{A}'_{j_1} \hookrightarrow \mathfrak{A}'_{j_n}$.

Recall that $\mathfrak{A}'_{j_n} = (\mathfrak{A}_{j_n}, \bar{a}_{j_n})$ whereby the image of \mathfrak{A}'_{j_1} under ι is a structure $(\mathfrak{B}, \bar{a}_{j_n})$. Then \mathfrak{B} has the following properties: (i) $\mathfrak{B} \in \mathcal{S}$, since $\mathfrak{A}_{j_1} \in \mathcal{S}$ and \mathcal{S} is closed under isomorphisms, (ii) $\mathfrak{B} \subseteq \mathfrak{A}_{j_n}$, (iii) $\bar{a}_{j_n} \in \mathcal{U}_{\mathfrak{B}}^k$, (iv) $|\mathfrak{B}| = |\mathfrak{A}'_{j_1}| = r \leq j_n$, and (v) $(\mathfrak{B}, \bar{a}_{j_n}) \equiv_{m, \mathcal{L}} (\mathfrak{A}_{j_n}, \bar{a}_{j_n})$. This contradicts the property of \mathfrak{A}_{j_n} stated at the outset, completing the proof. \square

Remark 11.2.3. The implication given by Theorem 11.2.2 does not in general, imply the existence of a computable witness function for \mathcal{L} -EBSP(\mathcal{S}, k). Consider the class \mathcal{S} of two dimensional grid posets; \mathcal{S} can be seen to be w.q.o. under embedding, whereby \mathcal{L} -EBSP($\mathcal{S}, 0$) holds. But if there is a computable witness function for \mathcal{L} -EBSP($\mathcal{S}, 0$), then by Lemma 11.1.1, it follows that \mathcal{L} -Th(\mathcal{S}) is decidable. Equivalently, the satisfiability problem for \mathcal{L} (the problem of deciding if a given \mathcal{L} sentence is satisfiable) is decidable over \mathcal{S} , for both $\mathcal{L} = \text{FO}$ and $\mathcal{L} = \text{MSO}$. However, this contradicts the known result that the MSO satisfiability is undecidable over two dimensional grid posets. The latter class of posets is thus an example of a class of structures that is w.q.o. under embedding, and hence satisfies MSO-EBSP($\cdot, 0$), but for which there is no computable witness function for MSO-EBSP($\cdot, 0$).

We now show that the converse to Theorem 11.2.2 does not hold in general.

Proposition 11.2.4. *There exists a class \mathcal{S} of structures such that FO-EBSP($\mathcal{S}, 0$) holds but WQO($\mathcal{S}, 0$) fails.*

Proof. Let C_n (respectively, P_n) denote an undirected cycle (respectively, path) of length n . Let mP_n denote the disjoint union of m copies of P_n . Let $H_n = \bigsqcup_{i=0}^{i=3^n} nP_i$ and $G_n = C_{3^n} \sqcup H_n$, where \sqcup denotes disjoint union. Now consider the class \mathcal{S} of undirected graphs that is closed under isomorphisms, and is given upto isomorphisms by $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2$, where $\mathcal{S}_1 = \{H_n \mid n \geq 1\}$ and $\mathcal{S}_2 = \{G_n \mid n \geq 1\}$. That WQO($\mathcal{S}, 0$) fails is easily seen by considering the sequence $(G_n)_{n \geq 1}$, and noting that C_{3^n} cannot embed in C_{3^m} unless $m = n$. We now show FO-EBSP($\mathcal{S}, 0$) holds with the witness function $\theta_{(\mathcal{S}, 0, \text{FO})}$ being given by $\theta_{(\mathcal{S}, 0, \text{FO})}(m) = |G_m|$. In other words, we show that for $\mathfrak{A} \in \mathcal{S}$ and $m \in \mathbb{N}$, there exists \mathfrak{B} such that (i) $\mathfrak{B} \in \mathcal{S}$, (ii) $\mathfrak{B} \subseteq \mathfrak{A}$, (iii) $|\mathfrak{B}| \leq \theta_{(\mathcal{S}, 0, \text{FO})}(m)$ and (iv) $\mathfrak{B} \equiv_m \mathfrak{A}$. Towards this, we first present some basic facts about C_n , P_n , H_n and G_n , that are easy to verify. Let $m \in \mathbb{N}$ be given.

(F.1) If $n_1, n_2 \geq 3^m$, then $P_{n_1} \equiv_m P_{n_2}$.

(F.2) If $n_1 \geq 3^m$ and $n_2 \geq m$, then $n_2 P_{n_1} \equiv_m m P_{3^m}$.

(F.3) If $n_1 \leq n_2$, then H_{n_1} always embeds in H_{n_2} .

(F.4) If $m \leq n_1 \leq n_2$, then $H_{n_1} \equiv_m H_{n_2}$. (follows from (1) and (2) above)

(F.5) If $n \geq m$, then $G_n \equiv_m H_n$.

Consider a structure $\mathfrak{A} \in \mathcal{S}$ and let $m \in \mathbb{N}$. We have two cases: (a) $\mathfrak{A} \in \mathcal{S}_1$ (b) $\mathfrak{A} \in \mathcal{S}_2$.

$\mathfrak{A} \in \mathcal{S}_1$: Then $\mathfrak{A} = H_n$ for some n . If $n \leq m$, then taking \mathfrak{B} to be \mathfrak{A} , we see that FO-EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, 0, m, \text{null}, \theta_{(\mathcal{S}, 0, \text{FO})}$) is true, where null is the empty tuple. Else $n > m$. Then consider H_m . From the facts F.3 and F.4, we have that H_m embeds in H_n and that $H_m \equiv_m H_n$. Then taking \mathfrak{B} to be the isomorphic copy of H_m that is a substructure of H_n , we see that FO-EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, 0, m, \text{null}, \theta_{(\mathcal{S}, 0, \text{FO})}$) is indeed true.

$\mathfrak{A} \in \mathcal{S}_2$: Then $\mathfrak{A} = G_n$ for some n . If $n \leq m$, then taking \mathfrak{B} to be \mathfrak{A} , we see that FO-EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, 0, m, \text{null}, \theta_{(\mathcal{S}, 0, \text{FO})}$) is true. Else $n > m$. Then consider H_m . From the facts F.3, F.4 and F.5, we see that H_m embeds in G_n and that $H_m \equiv_m G_n$. Then taking \mathfrak{B} to be the isomorphic copy of H_m that is a substructure of G_n , we see that FO-EBSP-condition($\mathcal{S}, \mathfrak{A}, \mathfrak{B}, 0, m, \text{null}, \theta_{(\mathcal{S}, 0, \text{FO})}$) is indeed true. \square

Using Theorem 11.2.2 as a technique to show \mathcal{L} -EBSP(\cdot, k) for classes of structures

Let \mathcal{S} be the class of all n -dimensional grid posets (i.e. cartesian product of linear orders, iterated n times), for a given $n \in \mathbb{N}$. From Theorem 10.2.2, it follows that FO-EBSP(\cdot, k) holds of the class of all linear orders (and with a computable witness function). Then using Lemma 10.4.4 and Theorem 10.4.5, we see that FO-EBSP(\mathcal{S}, k) is true for all $k \in \mathbb{N}$ (and with a computable witness function). But these results do not tell us whether MSO-EBSP(\mathcal{S}, k) is true. We demonstrate below that we can use Theorem 11.2.2 to show that MSO-EBSP(\mathcal{S}, k) is indeed true. Thus Theorem 11.2.2 gives us a new technique to show \mathcal{L} -EBSP(\cdot, k) for classes of structures for which the \mathcal{L} -EBSP(\cdot, k) property cannot be inferred (at least prima facie) using the results presented in Chapter 10.

We show that MSO-EBSP(\mathcal{S}, k) holds by showing that WQO(\mathcal{S}, k) holds. We show the latter for the case when \mathcal{S} is the class of all 2-dimensional grid posets. The proof for the case of r -dimensional grid posets for $r > 2$ can be done similarly.

Consider an infinite sequence $(G_i, \bar{a}_i)_{i \geq 0}$ of structures of \mathcal{S}^k , where G_i is a 2-dimensional grid poset and \bar{a}_i is a k -tuple from G_i , for $i \geq 1$. Let G_i be the cartesian product of linear orders $L_{i,1}$ and $L_{i,2}$, and let \bar{b}_i and \bar{c}_i be the projections of \bar{a}_i onto $L_{i,1}$ and $L_{i,2}$ respectively (in other words, \bar{b}_i is the k -tuple of first components of the elements of \bar{a}_i , while \bar{c}_i is the k -tuple of the

second components of the elements of \bar{a}_i). Now $(L_{i,1}, \bar{b}_i)$ can be looked at as a word $w_{i,1}$ over the powerset of $\{1, \dots, k\}$, such that (i) the underlying linear order of $w_{i,1}$ is $L_{i,1}$, and (ii) each position e of $w_{i,1}$ is labeled with the set of all those indices r in $\{1, \dots, k\}$ such that e equals the r^{th} component of \bar{b}_i . Similarly $(L_{i,2}, \bar{c}_i)$ can be looked at as a word $w_{i,2}$. Let M_i be the cartesian product of the words $w_{i,1}$ and $w_{i,2}$, and let N_i be labeled grid poset obtained from M_i such that (i) the unlabeled grid underlying N_i is exactly the same as the unlabeled grid underlying M_i (and the latter is the same as G_i), and (ii) the label of any element (g_1, g_2) of N_i is the intersection of the labels of g_1 and g_2 in $w_{i,1}$ and $w_{i,2}$. It is easy to see that N_i is simply a “coloured” representation of (G_i, \bar{a}_i) . Whereby if for $i, j \geq 0$, we have $N_i \hookrightarrow N_j$, then $(G_i, \bar{a}_i) \hookrightarrow (G_j, \bar{a}_j)$. The proof that $\text{WQO}(\mathcal{S}, k)$ holds is therefore completed by showing that indeed there exist $i, j \geq 0$ such that $i < j$ and $N_i \hookrightarrow N_j$.

Consider the sequences $(w_{i,1})_{i \geq 0}$ and $(w_{i,2})_{i \geq 0}$. Since words are w.q.o. under embedding (Higman’s lemma) and the cartesian product of two w.q.o. sets is also w.q.o. under the point-wise order, there exist i, j such that $i < j$ and the pair $(w_{i,1}, w_{i,2}) \hookrightarrow (w_{j,1}, w_{j,2})$ where \hookrightarrow for pairs means point-wise \hookrightarrow . Then $w_{i,1} \hookrightarrow w_{j,1}$ and $w_{i,2} \hookrightarrow w_{j,2}$, whereby $M_i \hookrightarrow M_j$, and hence $N_i \hookrightarrow N_j$.

On a final note for this section, we observe that the implication given by Theorem 11.2.2, taken in its contrapositive form, gives a *logic-based tool* to show non-w.q.o.-ness of a class of structures under isomorphic embedding.

11.3 \mathcal{L} -EBSP(\cdot, k) and the homomorphism preservation theorem

The homomorphism preservation theorem (HPT) is one of the important classical preservation theorems that has been of significant interest in the finite model theory setting [6, 16, 60, 61]. While the theorem was shown to be true over various special classes of finite structures (such as those seen earlier in Chapter 8, namely classes that are acyclic, of bounded degree or of bounded tree-width [6]), its status over the class of all finite structures was open for a long time. In a landmark paper [70], Rossman proved that this theorem is indeed a rare classical preservation theorem that holds over the class of all finite structures. However, Rossman’s result does not imply anything about the truth of the HPT over the aforementioned special classes of finite

structures, since restricting the theorem to special classes weakens both the hypothesis and the conclusion of the theorem. In this section, we show that the homomorphism preservation theorem, in fact a parameterized generalization of it along the lines of $\text{GLT}(k)$, holds over classes that satisfy \mathcal{L} -EBSP(\cdot, k).

We first formally define the notion of homomorphism and state the HPT. While the HPT holds for arbitrary vocabularies, we restrict our discussion to vocabularies τ containing only relation symbols. Given a vocabulary τ and τ -structures \mathfrak{A} and \mathfrak{B} , a *homomorphism* from \mathfrak{A} to \mathfrak{B} , denoted $h : \mathfrak{A} \rightarrow \mathfrak{B}$, is a function $h : U_{\mathfrak{A}} \rightarrow U_{\mathfrak{B}}$ such that $(a_1, \dots, a_n) \in R^{\mathfrak{A}}$ implies $(h(a_1), \dots, h(a_n)) \in R^{\mathfrak{B}}$ for every n -ary relation symbol $R \in \tau$. We say that an FO sentence φ is *preserved under homomorphisms* over a class \mathcal{S} of structures if for all structures $\mathfrak{A}, \mathfrak{B} \in \mathcal{S}$, if $\mathfrak{A} \models \varphi$ and there is a homomorphism from \mathfrak{A} to \mathfrak{B} , then $\mathfrak{B} \models \varphi$. We say an FO formula is *existential-positive* if it is built up from un-negated atomic formulas using conjunction, disjunction and existential quantification. The HPT characterizes preservation under homomorphisms using existential-positive sentences.

Theorem 11.3.1 (HPT). *A first order sentence is preserved under homomorphisms over all structures iff it is equivalent over all structures to an existential-positive sentence.*

We now define a parameterized generalization of the notion of preservation under homomorphisms, along the lines of preservation under k -ary covered extensions seen in Section 3.2. For this, we first define the notion of *k -ary homomorphic covering* as a parameterized generalization of the notion of homomorphism. Recall, for a vocabulary τ , that τ_k is the vocabulary obtained by expanding τ with k fresh and distinct constants c_1, \dots, c_k .

Definition 11.3.2. Let \mathfrak{A} be a τ -structure, and $k \in \mathbb{N}$. Let $\bar{a}_1, \dots, \bar{a}_t$ be an enumeration of the k -tuples of \mathfrak{A} , and let $I = \{1, \dots, t\}$. Let $R = \{(\mathfrak{B}_i, \bar{b}_i) \mid i \in I\}$ be a (non-empty) set of τ_k -structures. A *k -ary homomorphic covering from R to \mathfrak{A}* is a set \mathcal{H} of homomorphisms $\{h_i : (\mathfrak{B}_i, \bar{b}_i) \rightarrow (\mathfrak{A}, \bar{a}_i) \mid i \in I\}$. If \mathcal{H} exists, then we call R a *k -ary homomorphic cover* of \mathfrak{A} .

Remark 11.3.3. Observe that if $R = \{\mathfrak{B}\}$ for some τ -structure \mathfrak{B} , then R is a 0-ary homomorphic cover of \mathfrak{A} iff there is a homomorphism from \mathfrak{B} to \mathfrak{A} . Also for a structure \mathfrak{A} , if t is the number of k -tuples of elements of \mathfrak{A} , then for a set R of τ_k -structures, if there exists a k -ary homomorphic covering from R to \mathfrak{A} , then we require that $|R| = t$.

We now define the notion of *preservation under k -ary homomorphic coverings*. Recall from

Section 11.2 that for a class \mathcal{U} of τ -structures, \mathcal{U}^k denotes the class of all τ_k -structures whose τ -reducts are structures in \mathcal{U} .

Definition 11.3.4. Let \mathcal{S} be a class of structures and $k \in \mathbb{N}$. A subclass \mathcal{U} of \mathcal{S} is said to be *preserved under k -ary homomorphic coverings over \mathcal{S}* , abbreviated as \mathcal{U} is $h\text{-}PC(k)$ over \mathcal{S} , if for every collection R of structures of \mathcal{U}^k , if there is a k -ary homomorphic covering from R to \mathfrak{A} and $\mathfrak{A} \in \mathcal{S}$, then $\mathfrak{A} \in \mathcal{U}$. Given an \mathcal{L} -sentence ϕ , we say ϕ is $h\text{-}PC(k)$ over \mathcal{S} if the class of models of ϕ in \mathcal{S} is $h\text{-}PC(k)$ over \mathcal{S} .

A class of sentences that is $h\text{-}PC(k)$ over any class of structures is the class of, what we call, $(\forall^k \exists^*)$ -positive sentences. A formula φ is said to be $(\forall^k \exists^*)$ -positive if it is of the form $\forall x_1 \dots \forall x_k \psi(x_1, \dots, x_k)$ where $\psi(x_1, \dots, x_k)$ is an existential positive formula. Observe that for $k = 0$, the class of $(\forall^k \exists^*)$ -positive formulae is exactly the class of existential-positive formulae. We say that the *generalized HPT for \mathcal{L} and parameter k* , abbreviated $\mathcal{L}\text{-GHPT}(k)$, holds over a class \mathcal{S} if the following is true: An \mathcal{L} sentence ϕ is $h\text{-}PC(k)$ over \mathcal{S} iff ϕ is equivalent over \mathcal{S} to a $(\forall^k \exists^*)$ -positive (FO) sentence. Observe that $\text{FO-GHPT}(0)$ holds over a class \mathcal{S} iff HPT holds over \mathcal{S} . We show below that $\mathcal{L}\text{-GHPT}(k)$ holds over classes of structures that satisfy $\mathcal{L}\text{-EBSP}(\cdot, k)$. We in fact show something more general as we describe below. Towards this, we first present a “homomorphic” version of $\mathcal{L}\text{-EBSP}$.

Definition 11.3.5 ($h\text{-}\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$). Let \mathcal{S} be a class of finite structures and k be a natural number. We say that \mathcal{S} satisfies the *homomorphic \mathcal{L} -EBSP* for parameter k , abbreviated $h\text{-}\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ is *true*, if there exists a function $\theta_{(\mathcal{S}, k, \mathcal{L})} : \mathbb{N} \rightarrow \mathbb{N}$ such that for each $m \in \mathbb{N}$, for each structure \mathfrak{A} of \mathcal{S} and for every k -tuple \bar{a} from \mathfrak{A} , there exists $(\mathfrak{B}, \bar{b}) \in \mathcal{S}^k$ such that (i) there is a homomorphism $h : (\mathfrak{B}, \bar{b}) \rightarrow (\mathfrak{A}, \bar{a})$, (ii) $|\mathfrak{B}| \leq \theta_{(\mathcal{S}, k, \mathcal{L})}(m)$, and (iii) $\text{tp}_{\mathfrak{B}, \bar{b}, m, \mathcal{L}}(\bar{x}) = \text{tp}_{\mathfrak{A}, \bar{a}, m, \mathcal{L}}(\bar{x})$. We call $\theta_{(\mathcal{S}, k, \mathcal{L})}$ a *witness function* of $h\text{-}\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$.

The following lemma is easy to see and the proof is skipped.

Lemma 11.3.6. *Let \mathcal{S} be a class of structures. Then for each $k \in \mathbb{N}$, we have the following.*

1. $\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ implies $h\text{-}\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$
2. $h\text{-MSO-EBSP}(\mathcal{S}, k)$ implies $h\text{-FO-EBSP}(\mathcal{S}, k)$

Further, in each of the implications above, any witness function for the antecedent is also a witness function for the consequent.

We now state and prove the central result of this section.

Theorem 11.3.7. *Let \mathcal{S} be a class of finite structures and $k \in \mathbb{N}$ be such that $h\text{-}\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ is true. Then $\mathcal{L}\text{-GHPT}(k)$, and hence HPT, holds over \mathcal{S} . Further, if there is a computable witness function for $h\text{-}\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$, then the translation from an \mathcal{L} sentence that is $h\text{-PC}(k)$ over \mathcal{S} to an \mathcal{S} -equivalent $(\forall^k \exists^*)$ -positive sentence, is effective.*

The same statement as above holds when $h\text{-}\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ is replaced with $\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$.

Towards the proof of Theorem 11.3.7, we recall the notion of canonical conjunctive query from the literature. Given a τ -structure \mathfrak{A} of size n , the *canonical conjunctive query* associated with \mathfrak{A} , denoted $\xi_{\mathfrak{A}}$, is the sentence given by $\xi_{\mathfrak{A}} = \exists x_1 \dots \exists x_n \beta(x_1, \dots, x_n)$ where $\beta(x_1, \dots, x_n)$ is the conjunction of all atomic formulae of the form $R(x_{i_1}, \dots, x_{i_r})$ where $R \in \tau$, r is the arity of R , $i_1, \dots, i_r \in \{1, \dots, n\}$ and $(a_{i_1}, \dots, a_{i_r}) \in R^{\mathfrak{A}}$. Observe that $\xi_{\mathfrak{A}}$ is an existential-positive sentence. The following theorem by Chandra and Merlin characterizes when a homomorphism exists from a structure \mathfrak{A} to a structure \mathfrak{B} , in terms of $\xi_{\mathfrak{A}}$.

Theorem 11.3.8 (Chandra-Merlin, 1977). *Let \mathfrak{A} and \mathfrak{B} be two finite structures. Then there is a homomorphism from \mathfrak{A} to \mathfrak{B} iff $\mathfrak{B} \models \xi_{\mathfrak{A}}$.*

We now prove Theorem 11.3.7.

Proof of Theorem 11.3.7. Suppose $h\text{-}\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ is true.

‘If’ part of $\mathcal{L}\text{-GHPT}(k)$: Let ϕ be an \mathcal{L} sentence that is equivalent over \mathcal{S} to the $(\forall^k \exists^*)$ -positive sentence $\varphi = \forall^k \bar{x} \psi(\bar{x})$ where ψ is an existential-positive formula. Let $R = \{(\mathfrak{B}_i, \bar{b}_i) \in \mathcal{S}^k \mid i \in I\}$ be a set of structures from \mathcal{S}^k such that $\mathfrak{B}_i \models \phi$ for each $i \in I$. Let $\mathfrak{A} \in \mathcal{S}$ and suppose there exists a k -ary homomorphic covering \mathcal{H} from R to \mathfrak{A} . Consider a k -tuple \bar{a} from \mathfrak{A} . Since \mathcal{H} is a k -ary homomorphic covering, there exists $i \in I$ such that there is a homomorphism $h : (\mathfrak{B}_i, \bar{b}_i) \rightarrow (\mathfrak{A}, \bar{a}) \in \mathcal{H}$. Since $\mathfrak{B}_i \models \phi$, we have $\mathfrak{B}_i \models \varphi$ and hence $(\mathfrak{B}_i, \bar{b}_i) \models \psi(\bar{x})$. Since existential-positive formulas are preserved under homomorphisms, we have $(\mathfrak{A}, \bar{a}) \models \psi(\bar{x})$. Since \bar{a} is arbitrary, we have $\mathfrak{A} \models \varphi$, whence $\mathfrak{A} \models \phi$. Then ϕ is $h\text{-PC}(k)$ over \mathcal{S} .

‘Only if’ part of $\mathcal{L}\text{-GHPT}(k)$: Let ϕ be an \mathcal{L} sentence that is $h\text{-PC}(k)$ over \mathcal{S} . Let the rank of ϕ be m and let $p = \theta_{(\mathcal{S}, k, \mathcal{L})}(m)$, where $\theta_{(\mathcal{S}, k, \mathcal{L})}$ is a witness function of $h\text{-}\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$. Let $\text{Mod}(\mathcal{S}^k, \phi, p)$ be the set (upto isomorphism) of all models of ϕ in \mathcal{S}^k that have size $\leq p$. For $(\mathfrak{B}, \bar{b}) \in \text{Mod}(\mathcal{S}^k, \phi, p)$, let $\xi_{(\mathfrak{B}, \bar{b})}$ be the canonical conjunctive query associated with (\mathfrak{B}, \bar{b}) . Observe that $\xi_{(\mathfrak{B}, \bar{b})}$ is an $\text{FO}(\tau_k)$ sentence. Let $\xi_{(\mathfrak{B}, \bar{b})}[c_1 \mapsto x_1; \dots; c_k \mapsto x_k]$ be the formula whose free variables are among x_1, \dots, x_k , that is obtained by substituting

x_i for the free occurrences of c_i in $\xi_{(\mathfrak{B}, \bar{b})}$ for $i \in \{1, \dots, k\}$, where c_1, \dots, c_k are the constants of $\tau_k \setminus \tau$. We abuse notation slightly and denote $\xi_{(\mathfrak{B}, \bar{b})}[c_1 \mapsto x_1; \dots; c_k \mapsto x_k]$ simply as $\xi_{(\mathfrak{B}, \bar{b})}(x_1, \dots, x_k)$. Now consider the sentence $\varphi = \forall x_1 \dots \forall x_k \alpha(x_1, \dots, x_k)$ where $\alpha(x_1, \dots, x_k) = \bigvee_{(\mathfrak{B}, \bar{b}) \in \text{Mod}(\mathcal{S}^k, \phi, p)} \xi_{(\mathfrak{B}, \bar{b})}(x_1, \dots, x_k)$. Clearly φ is $(\forall^k \exists^*)$ -positive. We show below that ϕ is equivalent to φ over \mathcal{S} .

- $\phi \rightarrow \varphi$: Let $\mathfrak{A} \in \mathcal{S}$ be such that $\mathfrak{A} \models \phi$. Let $t = |\mathfrak{A}|^k$ and let $\bar{a}_1, \dots, \bar{a}_t$ be an enumeration of the k -tuples of \mathfrak{A} . Let $I = \{1, \dots, t\}$. Since $h\text{-}\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ is true, we have for each k -tuple \bar{a}_i from \mathfrak{A} where $i \in I$, that there exists a structure $(\mathfrak{B}_i, \bar{b}_i) \in \mathcal{S}^k$ such that (i) there is a homomorphism $h_i : (\mathfrak{B}_i, \bar{b}_i) \rightarrow (\mathfrak{A}, \bar{a}_i)$, (ii) $|\mathfrak{B}_i| \leq \theta_{(\mathcal{S}, k, \mathcal{L})}(m) = p$, and (iii) $\text{tp}_{\mathfrak{B}_i, \bar{b}_i, m, \mathcal{L}}(\bar{x}) = \text{tp}_{\mathfrak{A}, \bar{a}_i, m, \mathcal{L}}(\bar{x})$. Then $\mathfrak{B}_i \equiv_{m, \mathcal{L}} \mathfrak{A}$. Since the rank of ϕ is m , we have $\mathfrak{B}_i \models \phi$; then $(\mathfrak{B}_i, \bar{b}_i) \in \text{Mod}(\mathcal{S}^k, \phi, p)$. Let $\xi_{(\mathfrak{B}_i, \bar{b}_i)}$ be the canonical conjunctive query associated with $(\mathfrak{B}_i, \bar{b}_i)$, where $i \in I$. By Theorem 11.3.8, we have for each $i \in I$, that $(\mathfrak{A}, \bar{a}_i) \models \xi_{(\mathfrak{B}_i, \bar{b}_i)}$ whereby $(\mathfrak{A}, \bar{a}_i) \models \alpha(x_1, \dots, x_k)$. Then $\mathfrak{A} \models \varphi$.
- $\varphi \rightarrow \phi$: Let $\mathfrak{A} \in \mathcal{S}$ be such that $\mathfrak{A} \models \varphi$. As before, let $I = \{1, \dots, t\}$ and $\bar{a}_1, \dots, \bar{a}_t$ be an enumeration of the k -tuples of \mathfrak{A} . Since $\mathfrak{A} \models \varphi$, we have, recalling the form of φ , that for each $i \in I$, $(\mathfrak{A}, \bar{a}_i) \models \xi_{(\mathfrak{B}_i, \bar{b}_i)}(x_1, \dots, x_k)$ for some $(\mathfrak{B}_i, \bar{b}_i) \in \text{Mod}(\mathcal{S}^k, \phi, p)$. By Theorem 11.3.8, there is a homomorphism $h_i : (\mathfrak{B}_i, \bar{b}_i) \rightarrow (\mathfrak{A}, \bar{a}_i)$. Then the set $\{h_i \mid i \in I\}$ is a k -ary homomorphic covering from R to \mathfrak{A} . Since $(\mathfrak{B}_i, \bar{b}_i) \in \text{Mod}(\mathcal{S}^k, \phi, p)$, we have $(\mathfrak{B}_i, \bar{b}_i) \models \phi$ for each $i \in I$. Then since ϕ is $h\text{-PC}(k)$ over \mathcal{S} , we have $\mathfrak{A} \models \phi$.

That the above result holds when $h\text{-}\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ is replaced with $\mathcal{L}\text{-EBSP}(\mathcal{S}, k)$ follows directly from Lemma 11.3.6. \square

Chapter 12

Directions for future work

The results seen so far naturally motivate various questions that we propose as future work.

A. Questions regarding \mathcal{L} -EBSP(\cdot, k):

1. [Model-theoretic] The Łoś-Tarski theorem and the homomorphism preservation theorem are true over any class satisfying \mathcal{L} -EBSP(\cdot, k) (Theorems 9.1.2 and 11.3.7). What other theorems of classical model theory are true of classes satisfying \mathcal{L} -EBSP(\cdot, k)? For instance, are Lyndon's positivity theorem and Craig's interpolation theorem true of \mathcal{L} -EBSP(\cdot, k) classes?
2. [Poset-theoretic] Theorem 11.2.2 shows us that w.q.o. under embedding entails \mathcal{L} -EBSP($\cdot, 0$). The converse however is not true: Proposition 11.2.4 gives a class that is not w.q.o. under embedding but for which \mathcal{L} -EBSP($\cdot, 0$) holds. However, this class is not hereditary. This motivates the following question: Under what reasonable closure assumptions on a class \mathcal{S} does \mathcal{L} -EBSP($\mathcal{S}, 0$) become equivalent to w.q.o. under embedding? Another natural question, given Remark 11.2.3, is the following: what strengthening of the w.q.o. under embedding property entails \mathcal{L} -EBSP($\cdot, 0$) with *computable* witness functions?
3. [Relational structures whose Gaifman graphs are n -partite cographs] Given a τ -structure \mathfrak{A} where τ is relational, the *Gaifman graph* of \mathfrak{A} is an undirected graph $\mathcal{G}(\mathfrak{A}) = (V, E)$ such that V is exactly $U_{\mathfrak{A}}$, and for $a, b \in V$, the pair $(a, b) \in E$ iff for some r -ary relation $R \in \tau$ and some r -tuple $\bar{c} \in R^{\mathfrak{A}}$, it is the case that \bar{c} contains a and b as components. We can now ask whether our results showing \mathcal{L} -EBSP(\cdot, k) for graphs (cf. Theorem 10.3.1) can be lifted to relational structures via the Gaifman graphs of the latter. Specifically, is it the case under suitable assumptions, that a class of finite relational structures whose Gaifman graphs form a subclass of n -partite cographs, satisfies \mathcal{L} -EBSP(\cdot, k), and further, with a computable witness function? As a step in this direction, we indeed have been able to show that a

hereditary class of relational structures whose Gaifman graphs are of bounded tree-depth, satisfies $\text{FO-EBSP}(\cdot, k)$ with a computable witness function (see Theorem 4 of [75]).

4. [Computational] For what classes of structures that satisfy $\mathcal{L}\text{-EBSP}(\cdot, k)$, is it the case that there are elementary witness functions (as opposed to just computable witness functions)? For the case of words, trees and nested words, $\mathcal{L}\text{-EBSP}(\cdot, k)$ holds with necessarily non-elementary witness functions. This is because for any class \mathcal{S} satisfying $\mathcal{L}\text{-EBSP}(\cdot, k)$ with witness function $\theta_{(\mathcal{S}, k, \mathcal{L})}$, since any structure in the class is (m, \mathcal{L}) -similar to a structure of size $\leq \theta_{(\mathcal{S}, k, \mathcal{L})}(m)$, the index of the $\equiv_{m, \mathcal{L}}$ relation over \mathcal{S} is elementary if $\theta_{(\mathcal{S}, k, \mathcal{L})}$ is elementary. However over words, the index of the $\equiv_{m, \text{FO}}$ relation itself is non-elementary [27].

Since the model checking problem for MSO is fixed parameter tractable with elementary dependence on formula size, over classes of structures of bounded tree-depth or bounded shrub-depth, we would like to investigate if this elementariness shows up as the elementariness of the witness functions for the $\mathcal{L}\text{-EBSP}(\cdot, k)$ properties of the aforementioned classes. If so, this would also show that the index of the $\equiv_{m, \mathcal{L}}$ relation is elementary over these classes, as reasoned above.

5. [Concerning closure under operations] Is there a syntactic characterization of operations that are quantifier-free, monotone and $\equiv_{m, \mathcal{L}}$ -preserving? (cf. Theorem 10.4.11) Also, can Theorem 10.4.11 be generalized to $k > 0$?
6. [Structural] Is there a structural characterization of posets/graphs that satisfy $\mathcal{L}\text{-EBSP}(\cdot, k)$? If not in general, then under reasonable closure assumptions on the classes (like say hereditariness)? As a step in this direction, Theorem 8.2.2 shows that any hereditary class \mathcal{S} of directed graphs for which $\mathcal{L}\text{-EBSP}(\cdot, k)$ holds for any $k \geq 2$ (and hence over which $\text{GLT}(k)$ holds by Theorem 9.1.2) must be such that the underlying undirected graphs of the graphs of \mathcal{S} must have bounded induced path lengths. The converse of this statement is a technical challenging question, that we wish to investigate. Given the “empirical evidence” that many interesting classes of structures of interest in computer science satisfy $\mathcal{L}\text{-EBSP}(\cdot, \cdot)$, a structural characterization of the latter, even under reasonable assumptions (like hereditariness), might “give back” notions/new classes of structures of use and relevance to computer science. (As a very successful recent example of such a “give back”, a structural characterization under the assumption of hereditariness, of the notion of *quasi-wideness* that was introduced in [6] in the context of the homomorphism preservation theorem, yielded the notion of *nowhere dense graphs* [60, 61], and this class of graphs has turned out to be widely

useful from the combinatorial and algorithmic points of view [17, 22, 30].)

7. [Probabilistic] One can define a “probabilistic version” of \mathcal{L} -EBSP($\mathcal{S}, 0$) in which, instead of asserting that for any structure in \mathcal{S} , a bounded (m, \mathcal{L}) -similar substructure of it that is in \mathcal{S} , exists with probability 1, one asserts the same “with high probability”. One can define an analogous probabilistic version of \mathcal{L} -EBSP(\mathcal{S}, k). It would be interesting to investigate what classes of graphs satisfy this version of \mathcal{L} -EBSP(\cdot, k).

B. Questions regarding FO-GHPT(k):

1. Using techniques very similar to those presented in Section 4.1, and using *special models* (see Chp. 5 of [12]) instead of λ -saturated models, we can show the following result. Observe that FO-GHPT(0) is exactly HPT.

Theorem 12.1 (The generalized HPT). *Let \mathcal{S} be a class of arbitrary structures, that is elementary. Then FO-GHPT(k) holds over \mathcal{S} for each $k \in \mathbb{N}$.*

Over all finite structures, we know that FO-GHPT(0), which is HPT, is true by the results of Rossman [70]. Given that the HPT is amongst the very rare theorems from classical model theory to hold over all finite structures, it would be interesting to investigate if FO-GHPT(k) holds over all finite structures for $k > 0$.

C. Questions concerning GLT(k):

1. As Proposition 8.2.3 demonstrates, each of the classes of structures that are acyclic, or of bounded degree (more generally, wide), or of bounded tree-width fails to satisfy GLT(k) for $k \geq 2$. A natural question to investigate is the case of $k = 1$.
2. Proposition 8.1.1 shows for each $l \geq 0$, that $PSC(l)$ sentences cannot be equivalent to $\exists^k \forall^*$ sentences for any fixed $k \geq 0$. In particular, for each $k \geq 0$, Proposition 8.1.1 gives a sentence ψ_k that is PS , and hence $PSC(l)$ for each $l \geq 0$, over all finite structures, but that is not equivalent in the finite, to any $\exists^k \forall^*$ sentence. However, ψ_k is itself an $\exists^{k+1} \forall^*$ sentence, i.e. a Σ_2^0 sentence (cf. Remark 8.1.2).

This raises the following question: Is it the case that for each $l \geq 0$, any sentence that is $PSC(l)$ in the finite is equivalent in the finite, to a Σ_2^0 sentence? Recall that $PSC = \bigvee_{l \geq 0} PSC(l)$, and that every Σ_2^0 sentence is PSC over any class of structures. We can then reframe the aforesaid question as: Over all finite structures, is it the case that a sentence is PSC iff it is equivalent to a Σ_2^0 sentence? We conjecture that this is indeed the case.

Conjecture 12.2. *Over the class of all finite structures, a sentence is PSC if, and only if, it is equivalent to a Σ_2^0 sentence.*

Over arbitrary structures, PSC is characterized by Σ_2^0 as shown by Corollary 4.1.2. Then proving Conjecture 12.2 in the affirmative would give us a preservation theorem that is not only true over arbitrary structures but also true over all finite structures. It would be interesting to investigate (the relativized version of) this conjecture over the special classes of structures mentioned in the previous point, and also over the classes considered in the context of the homomorphism preservation theorem (such as nowhere dense classes).

Chapter 13

A summary of our contributions

We conclude by summarizing the contributions of this thesis in the classical and finite model theory settings. In each of these settings, our contributions are of three kinds: notions, results and techniques.

Classical model theory:

- A. Notions: We introduce the properties of *preservation under substructures modulo k -cruxes* ($PSC(k)$) and *preservation under k -ary conerved extensions* ($PCE(k)$) as natural parameterized generalizations of the classical properties of preservation under substructures and preservation under extensions (Definitions 3.1.1 and 3.2.4). Our properties are finitary and combinatorial, and are non-trivial both over arbitrary structures as well as over finite structures.
- B. Results:
 - (a) The generalized Łoś-Tarski theorem for sentences ($GLT(k)$): This result provides semantic characterizations of the $\exists^k\forall^*$ and $\forall^k\exists^*$ classes of sentences (Theorem 4.1.1). Whereby, we get *finer characterizations* of the Σ_2^0 and Π_2^0 fragments of FO sentences than those in the literature, which are via notions like unions of ascending chains, intersections of descending chains, Keisler’s 1-sandwiches, etc. *None* of the latter notions relates the *count* of quantifiers to any model-theoretic properties. As a consequence of $GLT(k)$, we obtain new semantic characterizations of the Σ_2^0 and Π_2^0 classes of FO sentences (Corollary 4.1.2).
 - (b) New semantic characterizations of the Σ_2^0 and Π_2^0 classes of FO theories: These characterizations are obtained via “infinitary” variants of $PSC(k)$ and $PCE(k)$, namely, the notions of preservation under substructures modulo λ -cruxes and preservation under λ -

ary covered extensions respectively, for infinite cardinals $\lambda > \aleph_0$ (Theorems 5.2.1(1) and 5.1.1(2)).

- (c) Applications in proving inexpressibility results in FO: We give new and simple proofs of well-known inexpressibility results in FO, such as inexpressibility of acyclicity, connectedness, bipartiteness, etc., using our preservation theorems (Section 4.2.1).

C. Techniques: We introduce a novel technique of getting a syntactically defined FO theory equivalent to a given FO theory satisfying a semantic property, *by going outside of FO* (Lemma 5.2.15 and Proposition 5.2.16). The idea is to first express the semantic property in a syntactically defined fragment of an infinitary logic, and then use a “compiler-result” to translate the aforementioned infinitary sentences to equivalent FO theories, when these sentences are known to be equivalent to FO theories. The latter FO theories are obtained from suitable *finite approximations* of the infinitary sentences, that are defined syntactically in terms of the latter. We believe this technique of accessing the descriptive power of an infinitary logic followed by accessing the translation power of a compiler result, may have other applications.

Finite model theory:

A. Notions: We define a new logic based combinatorial property of finite structures that we call the \mathcal{L} -Equivalent Bounded Substructure Property \mathcal{L} -EBSP(\mathcal{S}, k) (Definition 9.1).

B. Results:

- (a) A strengthening of the classical result showing the failure of the Łoś-Tarski theorem in the finite: We show that there is a vocabulary τ such that for each k , there is an $\text{FO}(\tau)$ sentence that is preserved under substructures over the class \mathcal{S} of all finite structures but that is not equivalent over \mathcal{S} , to any $\exists^k \forall^*$ sentence (Theorem 8.1.1). The case of $k = 0$ of this result is the classical failure of the Łoś-Tarski theorem in the finite.
- (b) A preservation theorem that imposes structural restrictions: We show that under the assumption that a given class \mathcal{S} of graphs is hereditary, if $\text{GLT}(k)$ holds over \mathcal{S} , then \mathcal{S} must have bounded induced path lengths (Theorem 8.2.2).
- (c) Characterizing prenex FO sentences with two blocks of quantifiers: The preservation theorems studied over well-behaved classes, namely the Łoś-Tarski theorem and the homomorphism preservation theorem, characterize Σ_1^0 and Π_1^0 sentences – sentences that contain only one block of quantifiers – or subclasses of these. We characterize over various interesting classes of finite structures, Σ_2^0 and Π_2^0 sentences – sentences which

contain two blocks of quantifiers.

- (d) Strong connections of \mathcal{L} -EBSP(\mathcal{S}, k) with classical model theory: The property of \mathcal{L} -EBSP(\mathcal{S}, k) entails GLT(k) (and hence the Łoś-Tarski theorem) as well as a generalization of the homomorphism preservation theorem, and even “effective” versions of all these theorems if the witness function for \mathcal{L} -EBSP(\mathcal{S}, k) is computable (Theorems 9.1.2 and 11.3.7). Furthermore, from the very close resemblance of its definition to that of the downward Löwenheim-Skolem property, \mathcal{L} -EBSP(\mathcal{S}, k) can very well be regarded as a finitary analogue of the latter (Section 9.2). To the best of our knowledge, finitary analogues of intrinsically infinitary properties from classical model theory have rarely been studied earlier.
- (e) Strong connections of \mathcal{L} -EBSP(\mathcal{S}, k) with computer science: We show that a variety of classes of interest in computer science satisfy \mathcal{L} -EBSP(\cdot, k), and further, with computable witness functions. These include the classes of words, trees (unordered, ordered, or ranked), nested words, cographs, graphs of bounded tree-depth, graph classes of bounded shrub-depth and n -partite cographs (Theorems 10.2.2 and 10.3.1). We show that \mathcal{L} -EBSP(\cdot, \cdot) remains preserved under finite unions and finite intersections, and under taking subclasses that are hereditary or \mathcal{L} -definable (Lemma 10.4.1). Again, \mathcal{L} -EBSP(\cdot, \cdot) remains preserved under various well-studied operations from the literature that are implementable using quantifier-free translation schemes; these include unary operations like complementation, transpose and the line-graph operation, binary “sum-like” operations like disjoint union, join and pointed substitution, and binary “product-like” operations that include various kinds of products like cartesian, tensor, lexicographic and strong products (Corollary 10.4.7). While it follows that \mathcal{L} -EBSP(\cdot, \cdot) remains preserved under finite unions of classes obtained by finite compositions of the above operations, we show that \mathcal{L} -EBSP($\cdot, 0$) remains preserved even infinite unions of such classes, provided these unions are “regular” (Theorem 10.4.11). These various closure properties enables us to construct a wide spectrum of classes of finite structures that satisfy \mathcal{L} -EBSP(\cdot, \cdot), and that are hence “well-behaved” model-theoretically. All of these classes are different from the well-behaved classes considered in the literature [6, 7, 38], and were earlier not known to enjoy the many model-theoretic properties that they do.
- (f) New composition results for nested words and n -partite cographs: Composition results

allow inferring the formulas that are satisfied in a structure that is built up from smaller structures, from the formulas satisfied in the latter structures [57]. Composition results for FO and MSO have traditionally been known for words. These have natural extensions to (unordered, ordered and ranked) trees. (We prove these extensions in this thesis.) We provide new FO and MSO composition results for nested words and n -partite cographs by defining the operations of “insert” and “merge” for these classes respectively, and showing that these operations possess the FO and MSO composition properties (Lemmas 10.2.6 and 10.3.2).

- (g) A new connection between well-quasi-ordering and logic: We show that any class of structures that is well-quasi-ordered (w.q.o.) under embedding satisfies \mathcal{L} -EBSP($\cdot, 0$) (Theorem 11.2.2). In contrapositive form, this result gives a logic-based tool to show that a class of structures is not w.q.o. under embedding. This result also shows that classes that are w.q.o. under embedding satisfy the Łoś-Tarski preservation theorem. This fact does not seem to be well-known [32].

C. Techniques: We prove an abstract result concerning tree representations (Theorem 10.1.1), that takes as input a tree-representation of a structure and produces as output, a small subtree that represents a small and logically similar substructure of the original structure. The output structure is obtained by iteratively performing appropriate “prunings” of, and “graftings” within, the input tree representation, in a manner that preserves the substructure and “ (m, \mathcal{L}) -similarity” relations between the structures represented by the trees before and after the pruning and grafting operations. Two key technical elements that are employed to perform the aforementioned operations are the finiteness of the index of the “ (m, \mathcal{L}) -similarity” relation and the *type-transfer property* of the tree-representation. We utilize our abstract result in showing the \mathcal{L} -EBSP(\cdot, \cdot) property for the variety of classes of structures that we mentioned earlier. Given that many interesting classes of finite structures have natural representations using trees, it is possible that our abstract result has more applications than the ones indicated in this thesis.

Bibliography

- [1] Miklos Ajtai and Yuri Gurevich. Monotone versus positive. *J. ACM*, 34(4):1004–1015, October 1987. 4
- [2] Miklos Ajtai and Yuri Gurevich. Datalog vs first-order logic. *J. Comput. Sys. Sci.*, 49(3): 562 – 588, 1994. 4
- [3] Natasha Alechina and Yuri Gurevich. Syntax vs. semantics on finite structures. In *Structures in Logic and Computer Science. A Selection of Essays in Honor of A. Ehrenfeucht*, pages 14–33. Springer-Verlag, 1997. 4
- [4] Rajeev Alur and Parthasarathy Madhusudan. Visibly pushdown languages. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*, pages 202–211, 2004. 10
- [5] Rajeev Alur and Parthasarathy Madhusudan. Adding nesting structure to words. *J. ACM*, 56(3), 2009. 10, 101
- [6] Albert Atserias, Anuj Dawar, and Phokion G. Kolaitis. On preservation under homomorphisms and unions of conjunctive queries. *J. ACM*, 53(2):208–237, 2006. 4, 5, 30, 131, 138, 143
- [7] Albert Atserias, Anuj Dawar, and Martin Grohe. Preservation under extensions on well-behaved finite structures. *SIAM J. Comput.*, 38(4):1364–1381, 2008. 4, 79, 86, 143
- [8] Calixto Badesa. *The Birth of Model Theory: Löwenheim’s Theorem in the Frame of the Theory of Relatives*. Princeton University Press, 2004. 90
- [9] Achim Blumensath and Bruno Courcelle. On the monadic second-order transduction hierarchy. *Log. Meth. Comp. Sci.*, 6(2), 2010. 116

- [10] Egon Börger, Erich Grädel, and Yuri Gurevich. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer-Verlag, 1997. 6, 7
- [11] Andrea Calì, Georg Gottlob, and Andreas Pieris. New expressive languages for ontological query answering. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence, AAAI 2011, San Francisco, California, USA, August 7-11, 2011*, 2011. 7
- [12] Chen C. Chang and Howard J. Keisler. *Model Theory*. Elsevier Science Publishers, 3rd edition, 1990. 1, 2, 3, 17, 23, 24, 25, 26, 38, 139
- [13] Hubert Comon, Max Dauchet, Remi Gilleron, Christof Löding, Florent Jacquemard, Denis Lugiez, Sophie Tison, and Marc Tommasi. Tree automata techniques and applications. Available at: <http://www.grappa.univ-lille3.fr/tata>, 2007. release October 12, 2007. 101
- [14] Bruno Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990. 4
- [15] Anuj Dawar. On preservation theorems in finite model theory. In *Invited talk at the 6th Panhellenic Logic Symposium-Volos, Greece*, 2007. 4
- [16] Anuj Dawar. Homomorphism preservation on quasi-wide classes. *J. Comput. Sys. Sci.*, 76(5):324–332, 2010. 4, 131
- [17] Anuj Dawar and Stephan Kreutzer. Domination problems in nowhere-dense classes. In *Proceedings of the 29th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2009, IIT Kanpur, India, December 15-17, 2009*, pages 157–168, 2009. 139
- [18] Anuj Dawar, Martin Grohe, Stephan Kreutzer, and Nicole Schweikardt. Model theory makes formulas large. In *Proceedings of the International Colloquium on Automata, Languages and Programming, ICALP 2007, Wroclaw, Poland, July 9 – 13, 2007*, pages 913–924, 2007. 43
- [19] Reinhard Diestel. *Graph Theory*. Springer-Verlag, 2010. 2, 115, 118, 128
- [20] Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer-Verlag, 2013. 4, 7

- [21] David Duris. Extension preservation theorems on classes of acyclic finite structures. *SIAM J. Comput.*, 39(8):3670–3681, 2010. 4, 5, 86
- [22] Zdenek Dvorak, Daniel Král, and Robin Thomas. Deciding first-order properties for sparse graphs. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, Las Vegas, USA, October 23-26, 2010*, pages 133–142. IEEE Computer Society, 2010. 139
- [23] Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite Model Theory*. Springer-Verlag, 2005. 69
- [24] Michael Elberfeld, Martin Grohe, and Till Tantau. Where first-order and monadic second-order logic coincide. In *Proceedings of the 27th Annual IEEE Symposium on Logic in Computer Science, LICS 2012, Dubrovnik, Croatia, June 25-28, 2012*, pages 265–274, 2012. 11
- [25] Moshe Emmer, Zurab Khasidashvili, Konstantin Korovin, and Andrei Voronkov. Encoding industrial hardware verification problems into effectively propositional logic. In *Proceedings of Formal Methods in Computer Aided Design, FMCAD 2010, Lugano, Switzerland, October 20 - 23, 2010*, pages 137–144, 2010. 7
- [26] Ronald Fagin, Phokion G. Kolaitis, Renée J. Miller, and Lucian Popa. Data exchange: semantics and query answering. *Theor. Comput. Sci.*, 336(1):89–124, 2005. 7
- [27] Markus Frick and Martin Grohe. The complexity of first-order and monadic second-order logic revisited. *Ann. Pure Appl. Logic*, 130(1-3):3–31, 2004. 138
- [28] Jakub Gajarský and Petr Hliněný. Faster deciding MSO properties of trees of fixed height, and some consequences. In *Proceedings of the 32nd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2012, Hyderabad, India, December 15-17, 2012*, pages 112–123, 2012. 11
- [29] Jakub Gajarsky and Petr Hliněný. Kernelizing MSO properties of trees of fixed height, and some consequences. *Log. Meth. Comp. Sci.*, 11(19):1–26, 2015. 11
- [30] Jakub Gajarský, Petr Hliněný, Jan Obdržálek, Sebastian Ordyniak, Felix Reidl, Peter Rossmanith, Fernando S. Villaamil, and Somnath Sikdar. Kernelization using structural parameters on sparse graph classes. In Hans L. Bodlaender and Giuseppe F. Italiano,

- editors, *Proceedings of the 21st Annual European Symposium on Algorithms, ESA 2013, Sophia Antipolis, France, September 2-4, 2013*, pages 529–540. Springer-Verlag, 2013. 139
- [31] Robert Ganian, Petr Hlinený, Jaroslav Nešetřil, Jan Obdržálek, Patrice Ossona de Mendez, and Reshma Ramadurai. When trees grow low: Shrubs and fast MSO1. In *Proceedings of the 37th International Symposium on Mathematical Foundations of Computer Science, MFCS 2012, Bratislava, Slovakia, August 27-31, 2012*, pages 419–430, 2012. 11, 109, 111, 128
- [32] Jean Goubault-Larrecq and Sylvain Schmitz. Deciding piecewise testable separability for regular tree languages. working paper or preprint, February 2016. URL <https://hal.inria.fr/hal-01276119>. 144
- [33] Erich Grädel and Eric Rosen. On preservation theorems for two-variable logic. *Math. Log. Quart.*, 45:315–325, 1999. 4
- [34] Erich Grädel, Phokion G. Kolaitis, Leonid Libkin, Maarten Marx, Joel Spencer, Moshe Y. Vardi, Yde Venema, and Scott Weinstein. *Finite Model Theory and Its Applications*. Texts in Theoretical Computer Science. An EATCS Series. Springer-Verlag, 2005. 3, 69
- [35] Martin Grohe and Johann A. Makowsky. *Model Theoretic Methods in Finite Combinatorics*, volume 558 of *Contemporary Mathematics*. AMS, 2011. 11
- [36] Sumit Gulwani. Dimensions in program synthesis. In *Proceedings of the 12th International ACM SIGPLAN Symposium on Principles and Practice of Declarative Programming, PPDP '10*, pages 13–24. ACM, 2010. 7
- [37] Yuri Gurevich. Toward logic tailored for computational complexity. In Michael M. Richter et al., editor, *Computation and Proof Theory: Proceedings of the Logic Colloquium held in Aachen, July 1823, 1983, Part II*, page 175216. Springer-Verlag, 1984. 4, 79
- [38] Frederik Harwath, Lucas Heimberg, and Nicole Schweikardt. Preservation and decomposition theorems for bounded degree structures. *Log. Meth. Comp. Sci.*, 11(4), 2015. 4, 5, 86, 116, 143
- [39] Graham Higman. Ordering by Divisibility in Abstract Algebras. *Proc. London Math. Soc.*, s3-2(1):326–336, January 1952. 12

- [40] Wilfrid Hodges. *A Shorter Model Theory*. Cambridge University Press, 1997. 1, 2, 17, 23, 24
- [41] Wilfrid Hodges. Model Theory (Draft 20 Jul 00), 2000. URL <http://wilfridhodges.co.uk/history07.pdf>. 2
- [42] Wilfried Imrich and Sandi Klavzar. Recognizing hamming graphs in linear time and space. *Inf. Process. Lett.*, 63(2):91–95, 1997. 11
- [43] Mathew Joseph, Gabriel M. Kuper, and Luciano Serafini. Query answering over contextualized RDF knowledge with forall-existential bridge rules: Attaining decidability using acyclicity. In *Proceedings of the 29th Italian Conference on Computational Logic, CILC 2014, Torino, Italy, June 16-18, 2014.*, pages 210–224, 2014. 7
- [44] Mathew Joseph, Gabriel M. Kuper, Till Mossakowski, and Luciano Serafini. Query answering over contextualized RDF/OWL knowledge with forall-existential bridge rules: Decidable finite extension classes. *Semantic Web*, 7(1):25–61, 2016. 7
- [45] Howard J. Keisler. Theory of models with generalized atomic formulas. *J. Symb. Logic*, 25:1–26, 1960. 3, 48, 49, 56
- [46] Howard J. Keisler. Finite approximations of infinitely long formulas. In J. Addison, editor, *The Theory of Models*, pages 158–169. North-Holland Pub. Co., Amsterdam, 1965. 56, 58
- [47] Howard J. Keisler. *Model Theory for Infinitary Logic. Logic with Countable Conjunctions and Finite Quantifiers*, volume 62 of *Studies in Logic and the Foundations of Mathematics*. North-Holland Pub. Co., Amsterdam-London, 1971. 2, 64, 154
- [48] John L. Kelley. *General Topology*. Graduate Texts in Mathematics. Springer, 1975. 2
- [49] Markus Krötzsch and Sebastian Rudolph. Extending decidable existential rules by joining acyclicity and guardedness. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI 2011, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 963–968, 2011. 7
- [50] Joseph B. Kruskal. The theory of well-quasi-ordering: A frequently discovered concept. *J. Comb. Theory Ser. A*, 13(3):297 – 305, 1972. 12

- [51] Kenneth Kunen. *Set Theory: An Introduction to Independence Proofs*. Mathematical Programming Study. North-Holland Pub. Co., Amsterdam, 1980. 2
- [52] Michael Lampis. Model checking lower bounds for simple graphs. *Log. Meth. Comp. Sci.*, 10(1), 2014. 11
- [53] Maurizio Lenzerini. Data integration: A theoretical perspective. In *Proceedings of the 21st ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, PODS '02, pages 233–246. ACM, 2002. 7
- [54] Leonid Libkin. *Elements of Finite Model Theory*. Springer-Verlag, 2004. 3, 17, 69, 72, 73, 74
- [55] Per Lindström. A characterization of elementary logic. In Sören Halldén, editor, *Modality, Morality and Other Problems of Sense and Nonsense*, pages 189–191. Lund, Gleerup, 1973. 12, 23
- [56] Leopold Löwenheim. Über möglichkeiten im relativkalkül. *Mathematische Annalen*, 76 (4):447–470, 1915. 90
- [57] Johann A. Makowsky. Algorithmic uses of the Feferman-Vaught theorem. *Ann. Pure Appl. Logic*, 126(1-3):159–213, 2004. 11, 74, 102, 115, 144
- [58] Anatoly I. Maltsev. Untersuchungen aus dem Gebiete der mathematischen Logik. *Matematicheskii Sbornik*, n.s.(1):323–336, 1936. 90
- [59] Eric C. Milner. Basic WQO- and BQO-theory. In Ivan Rival, editor, *Graphs and Order: The Role of Graphs in the Theory of Ordered Sets and Its Applications*, pages 487–502. Springer-Verlag, 1985. 3, 12
- [60] Jaroslav Nešetřil and Patrice Ossona De Mendez. First order properties on nowhere dense structures. *J. Symb. Logic*, 75(03):868–887, 2010. 131, 138
- [61] Jaroslav Nešetřil and Patrice Ossona de Mendez. On nowhere dense graphs. *Eur. J. Comb.*, 32(4):600 – 617, 2011. 131, 138
- [62] Ruzica Piskac, Leonardo Mendonça de Moura, and Nikolaj Bjørner. Deciding effectively propositional logic using DPLL and substitution sets. *J. Autom. Reasoning*, 44(4):401–424, 2010. 7

- [63] Frank P. Ramsey. On a problem of formal logic. *Proc. London Math. Soc.*, s2-30:264–286, 1930. 6
- [64] Neil Robertson and Paul D. Seymour. Graph minors. V. Excluding a planar graph. *J. Comb. Theory Ser. B*, 41(1):92–114, 1986. 4
- [65] Neil Robertson and Paul D. Seymour. Graph Minors. XX. Wagner’s Conjecture. *J. Comb. Theory Ser. B*, 92(2):325–357, November 2004. 3
- [66] Eric Rosen. *Finite model theory and finite variable logics*. PhD thesis, University of Pennsylvania, 1995. 4
- [67] Eric Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6:427–439, 1997. 4
- [68] Eric Rosen. Some aspects of model theory and finite structures. *Bull. Symbolic Logic*, 8(3):380–403, 2002. 4, 79
- [69] Eric Rosen and Scott Weinstein. Preservation theorems in finite model theory. In *Logical and Computational Complexity. Selected Papers. Logic and Computational Complexity, International Workshop LCC ’94, Indianapolis, Indiana, USA, October 13-16, 1994*, pages 480–502, 1994. 4
- [70] Benjamin Rossman. Homomorphism preservation theorems. *J. ACM*, 55(3):15:1–15:53, 2008. 4, 30, 131, 139
- [71] Benjamin Rossman. *Personal Communication*. 2012. 43
- [72] Abhisekh Sankaran and Supratik Chakraborty. On semantic generalizations of the Bernays-Schönfinkel-Ramsey class with finite or co-finite spectra. *CoRR*, abs/1002.4334, 2010. 30
- [73] Abhisekh Sankaran, Bharat Adsul, Vivek Madan, Pritish Kamath, and Supratik Chakraborty. Preservation under substructures modulo bounded cores. In *Proceedings of the 19th International Workshop on Logic, Language, Information and Computation, WoLLIC 2012, Buenos Aires, Argentina, September 3-6, 2012*, pages 291–305, 2012. 7, 30, 155

- [74] Abhisekh Sankaran, Nutan Limaye, Akshay Sundararaman, and Supratik Chakraborty. Using preservation theorems for inexpressibility results in first order logic. Technical report, IIT Bombay, 2012. URL <http://www.cfdvs.iitb.ac.in/reports/index.php>. 42, 153
- [75] Abhisekh Sankaran, Bharat Adsul, and Supratik Chakraborty. A generalization of the Łoś-Tarski preservation theorem over classes of finite structures. In *Proceedings of the 39th International Symposium on Mathematical Foundations of Computer Science, MFCS 2014, Budapest, Hungary, August 25-29, 2014, Part I*, pages 474–485, 2014. 7, 138, 155
- [76] Abhisekh Sankaran, Bharat Adsul, and Supratik Chakraborty. A generalization of the Łoś-Tarski preservation theorem. *Ann. Pure Appl. Logic*, 167(3):189–210, 2016. 7, 155
- [77] Detlef Seese. Linear time computable problems and first-order descriptions. *Math. Struct. Comp. Sci.*, 6:505–526, 1996. 5
- [78] Thoralf Skolem. Logisch-kombinatorische untersuchungen über die erfüllbarkeit oder beweisbarkeit mathematischer sätze nebst einem theoreme über dichte mengen. *Videnskaps-selskapet Skrifter, I. Matematisk-naturvidenskabelig Klasse*, 4:1–36, 1920. 90
- [79] Saurabh Srivastava and Sumit Gulwani. Program verification using templates over predicate abstraction. In *Proceedings of the 30th Annual ACM SIGPLAN conference on Programming Language Design and Implementation, PLDI 2009, Dublin, Ireland, June 15-21, 2009*, pages 223–234, 2009. 7
- [80] Alexei P. Stolboushkin. Finitely monotone properties. In *Proceedings of the 10th Annual IEEE Symposium on Logic in Computer Science, LICS 1995, San Diego, USA, June 26-29, 1995*, pages 324–330. IEEE Computer Society, 1995. 4
- [81] William W. Tait. A counterexample to a conjecture of Scott and Suppes. *J. Symb. Logic*, 24(1):15–16, 1959. 4, 79
- [82] Alfred Tarski. The concept of truth in formalized languages. In A. Tarski, editor, *Logic, Semantics, Metamathematics*, pages 152–278. Oxford University Press, 1936. 1

Acknowledgments

I have a number of people to thank, who in their own unique ways have contributed to my growth during the long journey of my Ph.D., and an even longer stay in the IIT Bombay (IITB henceforth) campus. I have organized my acknowledgments in two parts: “academic” and “cultural”. In case I have missed out on thanking anyone, I apologize sincerely, and offer my sincere thanks to them at the outset itself.

A. “Academic” acknowledgments

First and foremost, I am indebted to my advisors Supratik Chakraborty and Bharat Adsul for the immense freedom and unconditional support they gave me to pursue my research for the length of time that I have. Indeed, but for these, I would never have been able to express myself through my research, to my satisfaction. For the many penetrating technical discussions, the valuable feedback on my results, writing and presentations, the sustained stipend (which I never had to bother about at all) and generous financial support for internships and conferences, and for all their appreciation of my work (which I feel is so vital for any creative activity), I offer my deepest gratitude to them.

I would like thank my panel members Ajit A. Diwan, Paritosh Pandya, and Akshay S., for providing me with their valuable comments over the years in my annual progress seminars. I thank Ajit A. Diwan additionally for various insightful discussions and for indeed being an “oracle” for my queries on graph theory and combinatorics. I thank Akshay S. again, along with Nutan Limaye, for always being helpful and appreciative, and for the discussions on inexpressibility results using preservation theorems [74].

It was my great fortune that I met Anand Pillay at the Logic and Set Theory satellite meeting of ICM (International Congress of Mathematicians) 2010, held at CMI (Chennai Mathematical Institute) in August 2010. I express my sincerest gratitude to him for having directed me to the subject of classical model theory for the characterization that I was after then ($\text{GLT}(k)$ indeed),

and for providing me, time and again, with the most valuable directions and advice, extended discussions and historical perspectives on the subject.

My heartfelt gratitude and thanks go to Anuj Dawar for the precious interactions and guidance that he provided me with on so many occasions on the subject of finite model theory. I thank him earnestly again for giving me the memorable opportunity of doing an internship with him during June-August 2012 at the University of Cambridge, UK. His deeply insightful observations and keen directions have proven to be of vital importance to many aspects of my thesis, particularly the finite model theory part of it.

I would like to thank the Association for Logic in India (ALI) for organizing the biennial Indian Conference on Logic and Applications (ICLA) and the biennial Indian School on Logic and Applications (ISLA) in alternate years. Likewise, I thank the Calcutta Logic Circle (CLC) for organizing the annual meets of CLC. Having been a regular at these meets for many years, I have benefitted a lot from them in terms of exposure to various aspects of logic, getting to meet a number of stalwarts, and getting platforms to present my work and obtain important feedback and opportunities. Members of ALI and CLC whom I would especially like to thank include Mohua Banerjee, Mihir Chakraborty, Soma Dutta, Sujata Ghosh, Kamal Lodaya, Ranjan Mukhopadhyaya, R. Ramajuman, Pulak Samanta, Jayanta Sen and Sourav Tarafder. I thank them again for all the warmth and touching affection that they have always given me. I thank R. Ramanajum further for many valuable discussions during the time I spent with him at the Institute of Mathematical Sciences (IMSc), Chennai in September-October 2010.

It is through the above mentioned meets that I got the privilege to meet and interact with Johan van Benthem and Rohit Parikh. I am very grateful to Johan van Benthem for sponsoring my first international trip to attend the European Summer School for Logic, Language and Information (ESSLLI) in July 2009 at Bordeaux, France. I am deeply thankful to Rohit Parikh for inviting me to spend a month at the Graduate Center, CUNY (City University of New York), USA in September-October 2012, and for so many memorable hours that we spent at New York and during his visits to Mumbai, discussing a variety of subjects, technical and otherwise. It was during my visit to CUNY that I had a providential meeting with Saharon Shelah. I offer him my very sincere thanks for appreciating $\text{GLT}(k)$, and for pointing me to [47] whose importance I realized only much later when contemplating on Theorem 5.2.3. I am very grateful to Thomas Scanlon for many valuable discussions on classical model theory, and sincerely thank him, Anand Pillay and Byunghan Kim, for inviting me to the ICM Satellite Conference on Clas-

sification Theory and its Applications at Daejeon, South Korea in August 2014, which exposed me to the frontier of one of the currently very active themes of classical model theory.

I am most thankful to Haim Gaifman, Erich Grädel, Martin Grohe, Yuri Gurevich, Lauri Hella, Phokion Kolaitis, Janos Makowsky, Jaroslav Nešetřil, Benjamin Rossman and Jouko Väänänen for giving me their valuable time, encouragement and discussions on finite model theory. I earnestly thank further, Erich Grädel for his warm hospitality during my visit to RWTH Aachen, Germany in September 2015, and Martin Grohe for raising the possibility that preservation under substructures may not be capturable by $\exists^k\forall^*$ sentences for any fixed k , over all finite structures – indeed I was able to confirm this later (Theorem 8.1.1). I express my special gratitude to Benjamin Rossman for the interactions we had at FSTTCS 2011 and ICM 2014, for sharing his unpublished result (Theorem 4.3.2), and for being such an inspiration for me ever since my Ph.D. took its turn to model theory.

The company of all of the above people is what I regard as, what is called in Sanskrit as, *satsang*, meaning the company of the wise and the elevated. I am grateful to all the people above for their *satsang*, and I look forward to continue to benefit from it in the future as well.

The international trips mentioned above, along with trips to other international conferences, were made possible due to financial supports from different sources. For their generous support, I gratefully acknowledge ACM India for sponsoring my travel to the ACM Turing Centenary Celebrations 2012 (USA), Indian Association for Research in Computing Sciences (IARCS) and IITB for my travel to MFCS 2012 (Hungary), and the CSE department of IITB for my travel to WoLLIC 2012 (Argentina), ICM 2014 (South Korea) and HLF 2015 (Germany).

In connection with my published papers, I thank the anonymous reviewers of [73, 75, 76] for their careful reading of the earlier versions of these, and for their many keen suggestions on improving the presentation of the results therein. I also thank Prithish Kamath and Vivek Madan, my collaborators in [73], for the many stimulating discussions we had as a part of their Bachelor's theses (and appreciate them for their mathematical maturity even as undergraduates) that led to more insights about $\text{GLT}(k)$, and also proving it in special cases (before it was shown to be true in general, over arbitrary structures, in [76]).

I would like to thank Amey Karkare for making the L^AT_EX thesis style file for IITB, that has been used in preparing this thesis. I also thank him sincerely for his technical help in preparing the two indices that can be found towards the end of this document. I am very grateful to the IITB CSE department for providing me a robust laptop to complete the preparation of the final

version of my thesis, at the critical juncture when my laptop went out of order just as I had begun to give my thesis its final edits.

I express many sincere thanks to the following people of the IITB CSE department for their affection, support and encouragement: Pushpak Bhattacharyya, Rushikesh Joshi, Saraswati Krithivasan, Krithi Ramamritham, Krishna S., Amitabha Sanyal, G. Sivakumar, S. Sudarshan and Sundar Viswanathan. I express my special gratitude to S. Sudarshan for his exceptional support during his tenure as the Head of the CSE Department (and for being, in my opinion, amongst the coolest HoDs the CSE department has ever had). Finally, I heartily thank the administrative staff of CFDVS (Center for Formal Design and Verification of Software), in particular Chandrakant Talekar, Kakasaheb Killedar and Pravin Jadhav, and the administrative staff of the CSE department, in particular Alpana Athavankar, Vijay Ambre, Sunanda Ghadge, Homcy, and Nitin Peje for completely taking care of all administrative matters in connection with my Ph.D., and further, always with so much warmth.

On a concluding note of my academic acknowledgments, I thank Prof. Amitabha Gupta and Mrs. Chaitali Gupta for always being so affectionate, and appreciative of my work. I hold them in high regard for their untiring efforts to make possible two surveys of modern logic (“Logic and Philosophy Today”, Vol. I and II published by College Publications) containing expository articles written by world experts on various aspects of logic, which I am sure will prove to be of immense value to students in India pursuing mathematics, computer science or philosophy.

B. “Cultural” acknowledgments

Music has been and will always be an inseparable part of my life. I express my deepest gratitude to all my teachers of Indian classical music (specifically, North Indian classical music or Hindustāni Shāstriya Sangeet). My sincerest pranām (a Sanskrit word that means bowing in reverence) to Shri Ajay Thosar for teaching me vocal music from absolute basics and with so much love always, for giving me the freedom to explore on my own right from day one, the ocean of Indian classical music, for being exceptionally considerate and patient with me, and for being a constant guide in my musical journey. My sincere pranām to Shri Shripad Zail from whom I took my initial lessons in tabla (a North Indian classical percussion instrument), for initiating me formally into the profoundly abstract world of Indian rhythm, so admirably and with such warmth and fondness. My sincere pranām to Shri Upendra Patnaik who took forward my tabla training, and so affectionately enriched me with several rhythmically sophisticated

compositions, as well as anecdotes from the lives of Indian classical musicians. Finally, it has been my great privilege to learn tabla under the tutelage of Pandit Nayan Ghosh. I am deeply indebted to him for imparting a wide repertoire of technically involved compositions and their corresponding playing techniques, a balanced understanding of both the rhythmic and melodic aspects of music, umpteen anecdotes from the lives of great masters, and for all the affection that he has given me; I offer him my most earnest pranām. Finally, my heartfelt pranām to the Spirit of Indian classical music for the enormous influence that it has had on my life and for the most unconditional benevolence that It has showered on me on so many occasions. To It, and to all the people above who show the way to It, I pledge my life-long studentship and service. My fervent thanks to my friends who made my stay at IITB so rich and full. I can never thank enough Meenarli, Prashant, Poornima, Vishal (Anvekar), Iman, Kasi and Russell for the committed friendship, whole-hearted affection, and unstinted support and care they have shown since the time I have known them. The limited space here restricts me to write so much more about each one of you. I would just say to you all – life has become so beautiful just because of you guys, and each one of you occupies a very very special and unique place in my heart.

I am grateful to my department seniors Hrishi, Ravi-g (or simply, Ravig), Shetal, Sobhan and Sreyash for all their affection (indeed, pampering), and words of wisdom they have given me over the years. I thank my department mates Anindya, Durgesh, Girish, Jagdish, Jinesh, Preeti, Rijurekha, Soumitra and Sujesha for being such wonderful people. My thanks to my lab mates for the awesome times we've had in the old and new labs. My special thanks and gratitude go to Rakesh for being ever present to help me out with problems arising out of my pedestrian understanding of technology, such as those pertaining to my laptop, mobile phone and useful services on the Internet (like Dropbox), and for the memorable times we spent doing creative design work during the organization of ALC 2015 and ICLA 2015 that took place at IITB.

Outside of the department, I would like to earnestly thank Dr. S. N. Jha, my football coach and former in-charge of the IITB Gymkhana, for his heart-warming care and affection, and personal attention that he always gave me, and for making me realize the crucial role of physical education in academics and research. I thank sincerely the members of the “Kannada gang” for the memorable times we have shared: Ajay, Karthik (SRK), Nanditha, Prasanna, Smita, Vineet, and Vinita. My thanks to my hostel mates Dev-Sir, Dushyant, Jai (JNT), Lalit, Mithun, Parkash, Pranab, Riju, Sathish and Vishvendra for making my stay in the hostel so much enjoyable and so lively. I thank my dearest canine friends Damroo, Kali, Kalia and Kalu for being so very loving

and caring, and for making life in the hostel such a celebration. I thank the people who fed me for all these years: my friends in the hostel mess (Bhagwat, Kailas, Munnabhai, Pandeyji, Rajesh and Shiva), in the hostel canteen (Dilip, Ganesh, Manoj, Sachin, Satish and Soma), at Staff Canteen (Hemant, Krishna Mama, Nilesh, Sanjay, Shankar Mama and Yogesh) and at Sunrise Dhaba (Mirza Uncle and Mukesh). The results in this thesis have a strong connection with the food that they so fondly prepared always. A special word of thanks to Mirza Uncle for his painstaking efforts in constructing the lovely ambience at Sunrise Dhaba, contemplations in which have led to insights for some results in this thesis (such as Theorem 5.2.1(1)).

As I near the end of my acknowledgments, I return to music since a majority of my friendships outside of the department, hostel and special “gangs” are those that music made possible. Surbahaar and Swarsandhya, the two annual music shows of IITB, have been an integral part of my life at IITB, and enabled me to bring about the much needed metamorphosis for my Ph.D. life from a very different undergraduate life. Of the people who were a part of it, there are so many to thank, but I would especially mention Amitji, Anubhav, Arpita (Mondal), Balu, Kolte, Sayantani, Shantanu, Srabasti, Sukhada and Suman. I look back at the time spent with such talented musicians and awesome people as you guys, with much nostalgia. Ujaan and Bijoya-Sammelani, the annual events of the Bengali Cultural Association of IITB, gave me an opportunity to know Madhumanti and Avradeep, two highly trained and proficient classical musicians, whom I thank for their great music and friendship. Performances at the cultural evenings of the conferences of various departments gave me an opportunity to know Hemali, and I thank her so much for being the wonderful friend that she is. Finally, I thank Milind Malshe for his Thursday *baithaks* (musical sessions), for generously sharing his deep insights into *rāgās* (Indian melodic scales) and for giving me opportunities to perform at his annual *gurupoornima* (an annual Indian festival dedicated to teachers) programmes.

On a concluding note, I thank my family members, cousins, aunts, uncles and grandparents, for their support and encouragement over the years. I especially mention my brother Kannan and sister-in-law Andrea, aunts Bhuvana Ramanathan and Subhadra Sethuraman, and grandmother Rajalakshmi Chandran for being constant sources of warmth and affection. Last but not the least, I express my indebtedness to my parents. I can never put into words what you mean for me. Your immeasurable sacrifices, unfathomable love, and being permanently there for me at all times, are things that I will never ever be able to repay. This thesis is dedicated to you.

Abhisekh Sankaran

Index to symbols

- (\mathfrak{A}, \bar{a}) , 23
 $(\mathfrak{A}_\eta)_{\eta < \lambda}$, 24
 PCE , 33
 $PCE(k)$, 32
 PSC , 30
 $PSC(\lambda), PCE(\lambda)$, 40
 $PSC(k)$, 29
 $\Delta_{\mathcal{L}}(m, \mathcal{S})$, 72
 $\Lambda_{\mathcal{S}, \mathcal{L}}$, 72
 Σ_n^0, Π_n^0 , 21
 Ξ , 75
 $\alpha, \beta, \gamma, \xi, \phi, \varphi, \chi, \psi$, 18, 70
 $\bigcup_{\eta < \lambda} \mathfrak{A}_\eta$, 24
 \mathcal{S}^k , 128
 \cong , 22
 $\text{Diag}(\mathfrak{A})$, 23
 $\text{El-diag}(\mathfrak{A})$, 23
 \equiv , 22
 $\exists^k \forall^*, \forall^k \exists^*$, 21
 $\exists^k \bar{x}, \exists \bar{x}, \forall^k \bar{x}, \forall \bar{x}, \exists^*, \forall^*$, 18
 $\text{Th}(\mathfrak{A})$, 25
 $\text{tp}_{\mathfrak{A}, \bar{a}}(x_1, \dots, x_k)$, 24
 $\text{GLT}(k)$, 35
 \hookrightarrow , 22
 $\lambda, \mu, \kappa, \eta$, 17
 \mathcal{L} -EBSP-condition, 87
 \mathcal{L} -EBSP(\mathcal{S}, k), 87
 $[\bigvee] \mathcal{F}$, 56
 $[\exists^k \wedge] \mathcal{F}$, 56
 \leftrightarrow , 21
 $\equiv_{m, \mathcal{L}}, \equiv_m$, 71, 72
 $\text{Th}_{m, \mathcal{L}}(\mathfrak{A})$, 71
 $\text{tp}_{\mathfrak{A}, \bar{a}, m, \mathcal{L}}(x_1, \dots, x_k)$, 71
 $\mathfrak{A}, \mathfrak{B}$, 19
 $\mathfrak{A}_{\mathfrak{B}}$, 23
 \models , 19, 20, 71
 $\text{tp}_{\Pi, \mathfrak{A}, \bar{a}}(x_1, \dots, x_k)$, 25
 \preceq , 22
 \mathcal{L} -RFSE, 99
 \subseteq , 22
 τ, σ , 17
 τ_λ , 22
 $\tau_{\mathfrak{B}}$, 23
 $\mathfrak{t} - \mathfrak{t}_{\geq a}$, 96
 $\mathfrak{t} [\mathfrak{t}_{\geq a} \mapsto \mathfrak{s}]$, 96
 $\mathfrak{t} \odot \mathfrak{s}$, 97
 $\mathfrak{t}_{\geq a}$, 96
 $\text{root}(\mathfrak{t})$, 96
 $\mathsf{U}_{\mathfrak{A}}, c^{\mathfrak{A}}, R^{\mathfrak{A}}, f^{\mathfrak{A}}, t^{\mathfrak{A}}$, 19
 \vdash , 21
 $\theta_{(\mathcal{S}, k, \mathcal{L})}$, 87

Index

- (m, \mathcal{L}) -type, 71
- Π_1^0 -type of a tuple, 25
- \mathcal{S} -equivalent, 71
- μ -saturation, 25
- τ -structure, 18
- k -ary cover, 32
- k -crux of a model, 29
- n -partite cograph, 109
- Łoś-Tarski theorem, 27

- ascending chain, 24
- atomic formula, 18, 69

- cartesian product, 77
- compactness theorem, 23
- composition lemma, 102
 - for n -partite cographs, 111
 - for nested words, 106
 - for ordered trees, 102
 - for unordered trees, 103
- consistent, 20

- diagram, 23
- disjoint union, 77
- downward Löwenheim-Skolem theorem, 24

- Ehrenfeucht-Fraïssé game, 72, 73

- elementarily
 - embeddable, 22
 - equivalent, 22
- elementary
 - chain theorem, 24
 - class of structures, 23
 - diagram, 23
 - embedding, 22
- embedding, 22
- entails, 21, 71
- equivalent, 21
- expansion
 - of a structure, 22
 - of a vocabulary, 22
- extension, 22
 - k -ary covered, 32
 - elementary, 22

- first order logic, 17
- formula, 18
- free variable, 18, 69, 70

- generalized Łoś-Tarski theorem, 35
- greatest common ancestor, 96

- hereditary, 2, 77

- homomorphism preservation theorem, 4, 131
- infinitary formulae, 56
 - finite approximations of, 56
- isomorphism, 22
- model, 19, 20, 71
- modulo a theory
 - $PSC(k)$, 29
 - entails, 21
 - equivalent, 21
- monadic second order logic, 69
- nested word, 101
- operation
 - tree, 123
 - implementable using quantifier-free translation scheme, 117
 - product-like, 115, 118
 - sum-like, 115, 118
- preservation under
 - λ -ary covered extensions, 40
 - k -ary covered extensions, 32
 - extensions, 25
 - substructures, 25
 - substructures modulo λ -cruxes, 40
 - substructures modulo k -cruxes, 29
- quantifier prefix, 21
- quantifier-free, 18, 70
- rank, 70
- reduct of a structure, 22
- reduction favourable
 - for degree, 98
 - for height, 97
 - size effective representation, 99
- representation
 - feasible, 97
 - map, 97
 - schema, 100
- sandwich, 48
- satisfiable, 20, 71
- sentence, 70
- substructure, 22
 - elementary, 22
 - modulo a k -crux, 29
- term, 17, 69
- theory, 20
- translation scheme, 75
 - dimension, 75
- tree
 - ordered, 96
 - ranked, 100
 - unordered, 95
- truth in a structure, 19, 20, 71
- tuple realizing a type, 24
- type determining a crux, 51, 57
- type of a tuple, 24
- union of a chain, 24
- universe, 19
- valid, 20
- vocabulary, 17
- well-quasi-order, 128
- witness, 31
- witness function, 87